arXiv:1404.7092v1 [cs.LO] 28 Apr 2014

Robustness against Power is PSpace-complete

Egor Derevenetc[1,2] and Roland Meyer[2]

[1]Fraunhofer ITWM [2]University of Kaiserslautern

Abstract. Power is a RISC architecture developed by IBM, Freescale, and several other companies and implemented in a series of POWER processors. The architecture features a relaxed memory model providing very weak guarantees with respect to the ordering and atomicity of memory accesses.
Due to these weaknesses, some programs that are correct under sequential consistency (SC) show undesirable effects when run under Power. We call these programs not robust against the Power memory model. Formally, a program is robust if every computation under Power has the same data and control dependencies as some SC computation.
Our contribution is a decision procedure for robustness of concurrent programs against the Power memory model. It is based on three ideas. First, we reformulate robustness in terms of the acyclicity of a happens-before relation. Second, we prove that among the computations with cyclic happens-before relation there is one in a certain normal form. Finally, we reduce the existence of such a normal-form computation to a language emptiness problem. Altogether, this yields a PSpace algorithm for checking robustness against Power. We complement it by a matching lower bound to show PSpace-completeness.

1 Introduction

To execute code as fast as possible, modern processors reorder operations. For example, Intel x86/x86-64 and SPARC processors implement the Total Store Ordering (TSO) memory model [13] which allows write buffering: store operations in each thread can be queued and get executed on memory later. Processors can also execute independent instructions out of program order as soon as the input data and computational units are available for them. This is an inherent feature of the POWER and ARM microprocessors [12]. Moreover, Power and ARM memory models, unlike TSO, do not guarantee store atomicity: one write can become visible to different threads at different times. They only ensure that all threads see stores to the same memory location in the same order; stores to different memory locations can be seen in different order by different threads.

All these optimizations are usually designed so that a single-threaded program has the illusion that its instructions are executed in program order. The picture changes in the presence of concurrency. Concurrent programs are often assumed to have sequentially consistent (SC) semantics [10]: each thread executes its operations in program order, stores become visible immediately to all threads. Concurrent programs may observe a difference from SC when run on

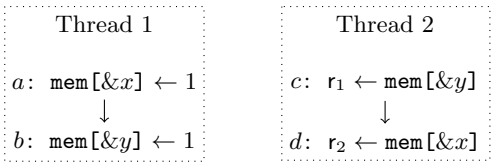

Fig. 1. Message Passing (MP) program [14]. By $\&x$ and $\&y$ we denote the addresses of the variables x and y. Initially, $x = y = 0$. The first thread writes a message into x and sets flag variable y, signifying that the message is written. The second thread reads the flag and, if it is set, expects to see the message written to x by the first thread.

a modern processor with a weak memory model. To see this, consider the MP program in Figure 1. SC and TSO forbid the situation where $\mathsf{r}_1 > \mathsf{r}_2$ upon termination of both threads. However, this is possible on Power: instruction c can read the value written by b, whereas d reads the initial value.

We call a program *not robust* against Power [15,6,7,2,5,8,4] if it exhibits non-SC behaviors when executed under the Power memory model. More formally, a program is robust if all its Power computations have the same data and control dependencies as the computations under SC. That is, for every Power computation there is a sequentially consistent computation which executes the same instructions, all loads read from the same stores in both computations, and stores to the same address happen in the same order. Robust programs produce the same results on Power and SC architectures, which means verification results for SC remain valid for the weak memory model.

We present an algorithm for deciding robustness against Power. This is the first decidability result for this architecture and, more generally, the first decidability result for a non-store atomic memory model. We obtain the algorithm in the following steps. First, we reformulate robustness in terms of acyclicity of a happens-before relation, using the result by Shasha and Snir [15]. Second, we show that among the computations with cyclic happens-before there is always one in a certain normal form. Next, we prove that the set of all normal-form computations can be generated by a multiheaded automaton — an automaton model developed recently in the context of robustness [8]. Finally, to check cyclicity of the happens-before relation we intersect this automaton with regular languages. The program is robust iff the intersection is empty. This reduces robustness to language emptiness for multiheaded automata. The algorithm works in space polynomial in the size of the program. We obtain a matching lower bound by a reduction of SC-reachability to robustness, similar to [5].

Related work The happens-before relation was formulated by Lamport [9]. Shasha and Snir [15] have shown that a computation violates sequential consistency iff it has a cyclic happens-before relation. Burckhardt and Musuvathi [6] proposed the first algorithm for detecting non-robustness against TSO based on monitoring SC computations. Burnim et al. [7] pointed out a mistake in the definition of TSO used in [6] and described monitoring algorithms for the TSO

and PSO memory models. Alglave and Maranget [2] presented a tool to statically over-approximate happens-before cycles in programs written in x86 and Power assembly, and to insert synchronization primitives (memory fences and syncs) as required for robustness (called stability in their work). Bouajjani et al. [5] obtained the first decidability result for robustness: robustness against TSO is PSPACE-complete for finite-state programs. In [4] they presented a reduction of robustness against TSO to SC reachability for general programs and an algorithm for optimal fence insertion.

The Power architecture has attracted considerable recent attention. Alglave et al. [3] give an overview of the numerous publications devoted to defining its semantics. We highlight two Power models: the operational model by Sarkar et al. [14] and the axiomatic one by Mador-Haim et al. [11]. These models were extensively tested against the architecture and were proven to be equivalent [11]. Nevertheless, the operational model is known to forbid certain behaviors that are possible on real hardware[1] and in the axiomatic model[2] [3]. Fortunately, there is a suggested fix: in Section 4.5 of [14] one should read *from a coherence-order-earlier write* instead of *from a different write* (two occurrences). Then, the operational model is believed to strictly and tightly over-approximate Power [1]. In the present paper we stick to the corrected operational model from [14].

Finally, we would like to note that ARM has a memory model very similar to that of Power. The differences and similarities are highlighted by Maranget et al. in [12,3]. This fact promises a relatively easy transfer of the proof techniques used in the present paper to the ARM memory model.

2 Programming Model

We define programs and their semantics in terms of automata. An *automaton* is a tuple $A = (S, \Sigma, \Delta, s_0, F)$, where S is a set of states, Σ is an alphabet, $\Delta \subseteq S \times (\Sigma \cup \{\varepsilon\}) \times S$ is a set of transitions, $s_0 \in S$ is an initial state, and $F \subseteq S$ is a set of final states. We call the automaton *finite* if S and Σ are finite. We write $s_1 \xrightarrow{a} s_2$ if $t = (s_1, a, s_2) \in \Delta$ and denote $\mathsf{src}(t) := s_1$, $\mathsf{dst}(t) := s_2$, $\mathsf{lab}(t) = a$. The *language* of the automaton is $\mathcal{L}(A) := \{\sigma \in \Sigma^* \mid s_0 \xrightarrow{\sigma} s \text{ for some } s \in F\}$. For a sequence $\sigma = a_1 \dots a_n \in \Sigma^*$ we define $|\sigma| := n$, $\sigma[i] := a_i$, $\mathsf{first}(\sigma) := a_1$, and $\mathsf{last}(\sigma) := a_n$. We use $\cdot$ for concatenation, $\downarrow$ for projection, and ε for the empty sequence. Given $\alpha \in \Sigma^*$ and $a, b \in \alpha$, we write $a <_\alpha b$ if $\alpha = \alpha_1 \cdot a \cdot \alpha_2 \cdot b \cdot \alpha_3$. Given a function $f\colon X \to Y$, $x' \in X$, and $y' \in Y$, we define $f' = f[x' \hookleftarrow y']$ by $f'(x) := f(x)$ for $x \in X \setminus \{x'\}$ and $f'(x') := y'$.

A program is a finite sequence of threads: $\mathcal{P} = \mathcal{T}_1 \dots \mathcal{T}_n$. A *thread* is an automaton $\mathcal{T}_{\mathsf{tid}} = (Q_{\mathsf{tid}}, \mathsf{CMD}, \mathcal{I}_{\mathsf{tid}}, q_{0\mathsf{tid}}, Q_{\mathsf{tid}})$ with a finite set of control states Q_{tid}, all of them being final, initial state $q_{0\mathsf{tid}}$, and a set of transitions $\mathcal{I}_{\mathsf{tid}}$ called *instructions* and labeled with *commands* CMD defined below. Each thread has an id from $\mathsf{TID} := [1..|\mathcal{P}|]$.

[1] `http://diy.inria.fr/cats/pldi-power/#lessvs`
[2] `http://diy.inria.fr/cats/cav-power/`

Let $\mathsf{DOM} = \mathsf{ADDR}$ be a finite domain of values and addresses containing the value 0. Let REG be a finite set of registers that take values from DOM. The set of commands CMD includes loads, stores, local assignments, and conditionals (`assume`):

$$\begin{aligned}\langle cmd \rangle \ ::= \ & \langle reg \rangle \leftarrow \texttt{mem[}\langle expr \rangle\texttt{]} \mid \texttt{mem[}\langle expr \rangle\texttt{]} \leftarrow \langle expr \rangle \\ \mid \ & \langle reg \rangle \leftarrow \langle expr \rangle \mid \texttt{assume(}\langle expr \rangle\texttt{)}\end{aligned}$$

The set of expressions EXPR is defined over constants from DOM, registers from REG, and (unspecified) functions FUN over $\mathsf{DOM} \cup \{\bot\}$. We assume that these functions return $\bot$ iff any of the arguments is $\bot$.

2.1 Power Semantics

We briefly recall the corrected model from [14]. The state of a running program consists of the runtime states of threads and the state of a storage subsystem.

The runtime state of a thread includes information about the instructions being executed by the thread. In order to start executing an instruction, the thread must *fetch* it. The thread can fetch any instruction whose source control state is equal to the destination state of the last fetched instruction. Then, the thread must perform any computation required by the semantics of this instruction. For example, for a load the thread must compute the address being accessed, then read the value at this address, and place it into the target register. The last step of executing an instruction is *committing* it. Committing an instruction requires committing all its *dependencies*. For example, before committing a load the thread must commit all its *address dependencies* — the instructions which define the values of registers used in the address expression — and *control dependencies* — the program-order-earlier (fetched earlier than the load) conditional instructions. Moreover, all loads and stores accessing the same address must be committed in the order in which they were fetched.

The storage subsystem keeps track, for each address, of the global ordering of stores to this address — the *coherence order* — and the last store to this address *propagated* to each thread. When a thread commits a store, this store is assigned a position in the coherence order which we identify by a rational number — the *coherence key*. We choose rational numbers (rather than naturals) to be able to insert a store between any two stores in the coherence order. The key must be greater than the coherence key of the last store to the same address propagated to this thread. The committed store is immediately propagated to its own thread. At some point later this store can be propagated to any other thread, as long as it is coherence-order-later (has a greater coherence key) than the last store to the same address propagated to that thread. When a thread loads a value from a certain address, it gets the value written by the last store to this address propagated to this thread. A thread can also forward the value being written by a not yet committed store to a later load reading the same address. This situation is called an *early read*.

An important property of Power is that it maintains the illusion of sequential consistency for single-threaded programs. This means that reorderings on

the thread level must not lead to situations when, e.g., a program-order-later load reads a coherence-order-earlier store than the one read by a program-order-earlier load from the same address. In [14] these restrictions are enforced by the mechanism of restarting operations. We put these conditions into the requirements on final states of the running program instead.

To keep the paper readable, we omit the descriptions of Power synchronization instructions: `sync`, `lwsync`, `isync`. All constructions in the paper can be consistently extended to support them with the final result continuing to hold.

Formally, we define the semantics of program $\mathcal{P}$ on Power by a *Power automaton* $Z(\mathcal{P}) := (S_Z, \mathsf{E}, \Delta_Z, s_{0Z}, F_Z)$. Here, E is a set of labels called *events* that we define together with the transitions.

State space A state of the Power automaton is a pair $s_Z = (\mathsf{ts}, s_Y) \in S_Z$ with runtime thread states $\mathsf{ts}\colon \mathsf{TID} \to S_X$ and storage subsystem state $s_Y \in S_Y$.

A runtime thread state $s_X = (\mathsf{fetched}, \mathsf{committed}, \mathsf{loaded}) \in S_X$ includes a finite sequence of fetched instructions $\mathsf{fetched} \in \mathcal{I}^*$, a set of indices of committed instructions $\mathsf{committed} \subseteq [1..|\mathsf{fetched}|]$, and a function giving the store read by a load $\mathsf{loaded}\colon [1..|\mathsf{fetched}|] \to \{\bot\} \cup \{\mathsf{init}_\mathsf{a} \mid \mathsf{a} \in \mathsf{ADDR}\} \cup \mathsf{TID} \times \mathbb{N}$. We use init_a to denote the initial store of value 0 to address a. The initial state of a running thread is $s_{0X} := (\varepsilon, \emptyset, \lambda i.\bot)$.

A state of the storage subsystem $s_Y = (\mathsf{co}, \mathsf{prop}) \in S_Y$ includes a mapping from a store instruction (its thread id and index in the list of fetched instructions) to its position in the coherence order $\mathsf{co}\colon (\mathsf{TID} \times \mathbb{N} \cup \{\mathsf{init}_\mathsf{a} \mid \mathsf{a} \in \mathsf{ADDR}\}) \to \mathbb{Q}$, and a mapping from a thread id and an address to the last store to this address propagated to this thread $\mathsf{prop}\colon \mathsf{TID} \times \mathsf{ADDR} \to \{\mathsf{init}_\mathsf{a} \mid \mathsf{a} \in \mathsf{ADDR}\} \cup \mathsf{TID} \times \mathbb{N}$. The initial state of the storage subsystem is $s_{0Y} := (\lambda\mathsf{tid}.\lambda i.0, \lambda\mathsf{tid}.\lambda\mathsf{a}.\mathsf{init}_\mathsf{a})$.

The initial state of automaton $Z(\mathcal{P})$ is $s_{0Z} := (\lambda\mathsf{tid}.s_{0X}, s_{0Y})$.

Transition relation Fix a state $s_Z = (\mathsf{ts}, s_Y)$ with $s_Y = (\mathsf{co}, \mathsf{prop})$ and a thread id $\mathsf{tid} \in \mathsf{TID}$ with runtime state $\mathsf{ts}(\mathsf{tid}) = (\mathsf{fetched}, \mathsf{committed}, \mathsf{loaded})$.

Let $\mathsf{eval}(\mathsf{tid}, i, e)$ return the value in DOM of expression e in the i'th fetched instruction of thread tid, or $\bot$ when the value is undefined. Formally $\mathsf{eval}(\mathsf{tid}, i, e) := \mathsf{v}$, where v is computed as follows. If $e \in \mathsf{DOM}$, then $\mathsf{v} := e$. If $e = \mathsf{f}(e_1 \dots e_n)$, then $\mathsf{v} := \mathsf{f}(\mathsf{eval}(\mathsf{tid}, i, e_1) \dots \mathsf{eval}(\mathsf{tid}, i, e_n))$. Otherwise, $e = \mathsf{r} \in \mathsf{REG}$. Let $i' \in [1..i-1]$ be the greatest index, such that $\mathsf{fetched}[i']$ is a local assignment or a load to r. If there is no such index, we define $\mathsf{v} := 0$. If $\mathsf{lab}(\mathsf{fetched}[i']) = \mathsf{r} \leftarrow e_\mathsf{v}$, then $\mathsf{v} := \mathsf{eval}(\mathsf{tid}, i', e_\mathsf{v})$. If $\mathsf{lab}(\mathsf{fetched}[i']) = \mathsf{r} \leftarrow \mathtt{mem}[e_\mathsf{a}]$, then $\mathsf{v} := \bot$ if $\mathsf{loaded}[i'] = \bot$, $\mathsf{v} := 0$ if $\mathsf{loaded}[i'] = \mathsf{init}_*$, and $\mathsf{v} := \mathsf{val}(\mathsf{loaded}[i'])$ otherwise (see the definition of val below).

The expression $\mathsf{addr}(\mathsf{tid}, i)$ returns the value of the address argument of the i'th fetched instruction of thread tid and is defined as follows. We use the special value $\top$ if the instruction has no such argument. If $\mathsf{lab}(\mathsf{fetched}[i]) = \mathsf{r} \leftarrow \mathtt{mem}[e_\mathsf{a}]$ or $\mathsf{lab}(\mathsf{fetched}[i]) = e_\mathsf{a} \leftarrow \mathtt{mem}[e_\mathsf{v}]$, then $\mathsf{addr}(\mathsf{tid}, i) := \mathsf{eval}(\mathsf{tid}, i, e_\mathsf{a})$. Otherwise, $\mathsf{addr}(\mathsf{tid}, i) := \top$.

Similarly, the expression $\mathsf{val}(\mathsf{tid}, i)$ returns the value of the value argument of the i'th fetched instruction of thread tid and is defined as follows. If $\mathsf{lab}(\mathsf{fetched}[i]) = \texttt{mem[}e_{\mathsf{a}}\texttt{]} \leftarrow e_{\mathsf{v}}$, $\mathsf{lab}(\mathsf{fetched}[i]) = \mathsf{r} \leftarrow e_{\mathsf{v}}$, or $\mathsf{lab}(\mathsf{fetched}[i]) = \texttt{assume(}e_{\mathsf{v}}\texttt{)}$, then $\mathsf{val}(\mathsf{tid}, i) = \mathsf{eval}(\mathsf{tid}, i, e_{\mathsf{v}})$. Otherwise, $\mathsf{val}(\mathsf{tid}, i) := \top$.

The expressions $\mathsf{addrdep}(\mathsf{tid}, i)$, $\mathsf{datadep}(\mathsf{tid}, i)$, $\mathsf{ctrldep}(\mathsf{tid}, i)$ denote the sets of indices of instructions in thread tid being respectively address, data, and control dependencies of the i'th instruction. The first two can be formally defined in a recursive manner, similar to eval. Also, $\mathsf{ctrldep}(\mathsf{tid}, i) := \{i' \in [1..i-1] \mid \mathsf{lab}(\mathsf{fetched}[i']) = \texttt{assume(}e_{\mathsf{v}}\texttt{)}\}$.

Let $\mathcal{T}_{\mathsf{tid}} = (Q_{\mathsf{tid}}, \mathsf{CMD}, \mathcal{I}_{\mathsf{tid}}, q_{0\mathsf{tid}}, Q_{\mathsf{tid}}) \in \mathcal{P}$. The transition relation Δ_Z is the smallest relation defined by the rules below:

POW-FETCH Consider $\mathsf{instr} \in \mathcal{I}_{\mathsf{tid}}$ with $\mathsf{src}(\mathsf{instr}) = \mathsf{dst}(\mathsf{last}(\mathsf{fetched}))$ or $\mathsf{src}(\mathsf{instr}) = q_{0\mathsf{tid}}$ if $\mathsf{fetched} = \varepsilon$, then:

$$(\mathsf{ts}, s_Y) \xrightarrow{(\mathsf{fetch},\mathsf{tid},\mathsf{instr})} (\mathsf{ts}[\mathsf{tid} \hookleftarrow (\mathsf{fetched} \cdot \mathsf{instr}, \mathsf{committed}, \mathsf{loaded})], s_Y).$$

POW-LOAD If $\mathsf{fetched}[i]$ is a load, $\mathsf{loaded}[i] = \bot$, $\mathsf{a} = \mathsf{addr}(\mathsf{tid}, i) \neq \bot$, then:

$$(\mathsf{ts}, s_Y) \xrightarrow{(\mathsf{load},\mathsf{tid},i,\mathsf{a})} (\mathsf{ts}[\mathsf{tid} \hookleftarrow (\mathsf{fetched}, \mathsf{committed}, \mathsf{loaded}[i \hookleftarrow \mathsf{prop}(\mathsf{tid}, \mathsf{a})])], s_Y).$$

POW-EARLY Let $\mathsf{fetched}[i]$ be a load, $\mathsf{loaded}[i] = \bot$, and $\mathsf{a} = \mathsf{addr}(\mathsf{tid}, i) \neq \bot$. Let $i' \in [1..i-1]$ be the greatest index such that $\mathsf{fetched}[i']$ is a store with $\mathsf{a}' = \mathsf{addr}(\mathsf{tid}, i') \in \{\mathsf{a}, \bot\}$. If $\mathsf{a}' \neq \bot$, $\mathsf{val}(\mathsf{tid}, i') \neq \bot$, $i' \notin \mathsf{committed}$, then:

$$(\mathsf{ts}, s_Y) \xrightarrow{(\mathsf{load},\mathsf{tid},i,\mathsf{a})} (\mathsf{ts}[\mathsf{tid} \hookleftarrow (\mathsf{fetched}, \mathsf{committed}, \mathsf{loaded}[i \hookleftarrow (\mathsf{tid}, i')])], s_Y).$$

POW-COMMIT Consider $i \in [1..|\mathsf{fetched}|] \setminus \mathsf{committed}$ where $\mathsf{fetched}[i]$ is not a store. Assume $\mathsf{addrdep}(\mathsf{tid}, i) \cup \mathsf{datadep}(\mathsf{tid}, i) \cup \mathsf{ctrldep}(\mathsf{tid}, i) \subseteq \mathsf{committed}$. Assume $\mathsf{a} = \mathsf{addr}(\mathsf{tid}, i) \neq \bot$ and $\mathsf{v} = \mathsf{val}(\mathsf{tid}, i) \neq \bot$. If $\mathsf{a} \neq \top$, assume $\{i' \in [1..i-1] \mid \mathsf{addr}(\mathsf{tid}, i') \in \{\mathsf{a}, \bot\}\} \subseteq \mathsf{committed}$. In case $\mathsf{fetched}[i]$ is a load, assume $\mathsf{loaded}[i] \neq \bot$. In case $\mathsf{fetched}[i]$ is an $\texttt{assume()}$, assume $\mathsf{v} \neq 0$. Then:

$$(\mathsf{ts}, s_Y) \xrightarrow{(\mathsf{commit},\mathsf{tid},i)} (\mathsf{ts}[\mathsf{tid} \hookleftarrow (\mathsf{fetched}, \mathsf{committed} \cup \{i\}, \mathsf{loaded})], s_Y).$$

POW-STORE Assume all the preconditions from the previous rule hold, but $\mathsf{fetched}[i]$ is a store. Choose a coherence key $\mathsf{k} \in \mathbb{Q}$ such that there is no $\mathsf{tid}' \in \mathsf{TID}$, $i' \in \mathbb{N}$ for which $\mathsf{co}(\mathsf{tid}', i') = \mathsf{k}$. Then:

$$(\mathsf{ts}, s_Y) \xrightarrow{(\mathsf{commit},\mathsf{tid},i,\mathsf{k},\mathsf{a})} (\mathsf{ts}[\mathsf{tid} \hookleftarrow (\mathsf{fetched}, \mathsf{committed} \cup \{i\}, \mathsf{loaded})], s'_Y),$$

where $s'_Y := (\mathsf{co}[(\mathsf{tid}, i) \hookleftarrow \mathsf{k}], \mathsf{prop})$.

Additionally, this transition is immediately followed by a POW-PROP transition propagating the store to the thread where it was committed.

POW-PROP Consider $\mathsf{tid}' \in \mathsf{TID}$, $i' \in \mathbb{N}$ with $\mathsf{co}(\mathsf{tid}', i') \neq \bot$. Let $\mathsf{a} = \mathsf{addr}(\mathsf{tid}', i')$. Assume $\mathsf{co}(\mathsf{prop}(\mathsf{tid}, \mathsf{a})) < \mathsf{co}(\mathsf{tid}', i')$. Then:

$$(\mathsf{ts}, s_Y) \xrightarrow{(\mathsf{prop},\mathsf{tid},\mathsf{tid}',i',\mathsf{a})} (\mathsf{ts}, (\mathsf{co}, \mathsf{prop}[(\mathsf{tid}, \mathsf{a}) \hookleftarrow (\mathsf{tid}', i')])).$$

Final states The set of final states $F_Z \subseteq S_Z$ consists of all states $s_Z = (\mathsf{ts}, (\mathsf{co}, \mathsf{prop})) \in S_Z$, such that for each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{ts}[\mathsf{tid}] = (\mathsf{fetched}, \mathsf{committed}, \mathsf{loaded})$ the following holds:

FIN-COMM All instructions are committed: $\mathsf{committed} = [1..|\mathsf{fetched}|]$.

FIN-LD Loads agree with the coherence order. Let $\mathsf{fetched}[i]$ be a load, and $\mathsf{fetched}[i']$ be an earlier load to the same address: $i' < i$, $\mathsf{addr}(\mathsf{tid}, i) = \mathsf{addr}(\mathsf{tid}, i')$. Then $\mathsf{co}(\mathsf{loaded}[i']) \leq \mathsf{co}(\mathsf{loaded}[i])$.

FIN-LD-ST Loads and stores in the same thread agree with the coherence order. Let $\mathsf{fetched}[i]$ be a load, let $\mathsf{fetched}[i']$ be an earlier store to the same address: $i' < i$, $\mathsf{addr}(\mathsf{tid}, i) = \mathsf{addr}(\mathsf{tid}, i')$. Then $\mathsf{co}(\mathsf{tid}, i') \leq \mathsf{co}(\mathsf{loaded}[i])$.

The set of all *Power computations of program* $\mathcal{P}$ is $\mathsf{C}_{\mathsf{power}}(\mathcal{P}) := \mathcal{L}(Z(\mathcal{P}))$. The set of all *SC computations of the program* $\mathsf{C}_{\mathsf{sc}}(\mathcal{P}) \subseteq \mathsf{C}_{\mathsf{power}}(\mathcal{P})$ includes only those computations where each instruction is executed atomically, and stores are immediately propagated to all threads.

Example 1. $\sigma_{MP} = \mathsf{fetch}(a) \cdot \mathsf{commit}(a) \cdot \mathsf{prop}(a,1) \cdot \mathsf{fetch}(b) \cdot \mathsf{commit}(b) \cdot \mathsf{prop}(b,1) \cdot \mathsf{prop}(b,2) \cdot \mathsf{fetch}(c) \cdot \mathsf{fetch}(d) \cdot \mathsf{load}(c) \cdot \mathsf{load}(d) \cdot \mathsf{commit}(d) \cdot \mathsf{commit}(c)$ is a feasible Power computation of the program MP (Figure 1):

- $\mathsf{fetch}(a) := (\mathsf{fetch}, 1, a)$ — thread 1 fetches store instruction a.
- $\mathsf{commit}(a) := (\mathsf{commit}, 1, 1, 1, \&x)$ — thread 1 commits a with $\mathsf{k} = 1$.
- $\mathsf{prop}(a, 1) := (\mathsf{prop}, 1, 1, 1, \&x)$ — a is propagated to its own thread.
- $\mathsf{fetch}(b) := (\mathsf{fetch}, 1, b)$ — thread 1 fetches store instruction b.
- $\mathsf{commit}(b) := (\mathsf{commit}, 1, 2, 2, \&y)$ — thread 1 commits b with $\mathsf{k} = 2$.
- $\mathsf{prop}(b, 1) := (\mathsf{prop}, 1, 1, 2, \&x)$ — the store is propagated to its thread.
- $\mathsf{prop}(b, 2) := (\mathsf{prop}, 2, 1, 2, \&x)$ — the store is propagated to thread 2.
- $\mathsf{fetch}(c) := (\mathsf{fetch}, 2, c)$ — thread 2 fetches load c.
- $\mathsf{fetch}(d) := (\mathsf{fetch}, 2, c)$ — thread 2 fetches load d.
- $\mathsf{load}(c) := (\mathsf{load}, 2, 1, \&y)$ — thread 2 reads value 1 written by b to y, because b was propagated to thread 2.
- $\mathsf{load}(d) := (\mathsf{load}, 2, 2, \&x)$ — thread 2 reads the initial value 0 of x, because a was not propagated to thread 2.
- $\mathsf{commit}(d) := (\mathsf{commit}, 2, 2)$ — thread 2 commits load d.
- $\mathsf{commit}(c) := (\mathsf{commit}, 2, 1)$ — thread 2 commits load c.

In the end, FIN-COMM holds as all fetched instructions are indeed committed, and FIN-LD and FIN-LD-ST trivially hold, as none of the threads has two instructions accessing the same address.

Lemma 1. *Assume* $s_{0Z} \xrightarrow{\sigma} s_Z \in F_Z$. *Then* s_Z *is uniquely determined.*

Proof. Given a state and an event e, there is at most one transition from this state labeled by e. This is clear for non-load events. For load events, this follows from Lemma 4 and Lemma 5: if a load event was produced by a load from memory transition, then condition (3) from Lemma 5 holds, then condition (1) from Lemma 4 cannot hold for any store, therefore, the load event cannot be produced by an early read transition. □

Lemma 2. *Let* $s_{0Z} \xrightarrow{\sigma} (\mathsf{ts}, s_Y) \xrightarrow{\mathsf{e}} (\mathsf{ts}', s_Y')$. *Let* $(\mathsf{fetched}, \mathsf{committed}, \mathsf{loaded}) = \mathsf{ts}(\mathsf{tid})$, $(\mathsf{fetched}', \mathsf{committed}', \mathsf{loaded}') = \mathsf{ts}'(\mathsf{tid})$ *for some* $\mathsf{tid} \in \mathsf{TID}$. *If* $\mathsf{loaded}[i] \neq \bot$, *then* $\mathsf{loaded}'[i] = \mathsf{loaded}[i]$.

Proof. Follows from the $\mathsf{loaded}[i] = \bot$ requirement in POW-LOAD and POW-EARLY transitions. □

Lemma 3. *Let* $s_{0Z} \xrightarrow{\sigma} s_Z \xrightarrow{\mathsf{e}} s_Z'$. *Assume* $\mathsf{eval}(\mathsf{tid}, i, e) = \mathsf{v} \neq \bot$ *in* s_Z. *Then* $\mathsf{eval}(\mathsf{tid}, i, e) = \mathsf{v}$ *in* s_Z'.

Proof. By definition of eval, Lemma 2, and the fact that functions in FUN are deterministic. □

Lemma 4. *Consider a computation* $\sigma \in \mathsf{C}_{\mathsf{power}}(\mathcal{P})$. *Then a load* (tid, i) *reads a value from a store* (tid, i') *via an early read (POW-EARLY) transition iff (1)* $\sigma = \sigma_1 \cdot (\mathsf{load}, \mathsf{tid}, i, \mathsf{a}) \cdot \sigma_2 \cdot (\mathsf{commit}, \mathsf{tid}, i', *, \mathsf{a}) \cdot \sigma_3$, $i' \in [1..i-1]$ *and (2)* σ_3 *does not contain events matching* $(\mathsf{commit}, \mathsf{tid}, [i'+1..i-1], *, \mathsf{a})$.

Proof. From left to right. Assume the load (tid, i) reads the store (tid, i') via an early read transition. Then (tid, i) must be the latest store to the same address in thread tid and must not be committed before load (i.e. committed after it), therefore (1) holds. If (2) does not hold, then (tid, i') is not the latest store to address a in thread tid before the load event, since stores to the same address are committed in the order of fetching. Contradiction.

From right to left. Let $s_{0Z} \xrightarrow{\sigma_1} s_Z = (\mathsf{ts}, s_Y)$. Consider $\mathsf{ts}(\mathsf{tid}) = (\mathsf{fetched}, \mathsf{committed}, \mathsf{loaded})$. Let $i'' < i$ be the greatest index, such that $\mathsf{fetched}[i'']$ is a store, $\mathsf{addr}(i'') \in \{\mathsf{a}, \bot\}$.

Assume $i' < i''$. If $\mathsf{addr}(i'') = \mathsf{a}$, we get a contradiction to (2), since stores to the same address are committed in the order of fetching. If $\mathsf{addr}(i'') = \bot$, then an early read is not possible in state s_Z, and the load reads from the latest propagated store (POW-LOAD), which is coherence-order-before the store (tid, i'), which is program-order-before (tid, i). This situation is forbidden by FIN-LD-ST.

By Lemma 3, $\mathsf{addr}(\mathsf{tid}, i') \in \{\mathsf{a}, \bot\}$, therefore, $i'' = i'$. Assume $\mathsf{addr}(\mathsf{tid}, i') = \bot$ or $\mathsf{val}(\mathsf{tid}, i') = \bot$. Then, again, a load from the latest propagated store takes place, which is impossible (see above). Therefore, $\mathsf{addr}(\mathsf{tid}, i') = \mathsf{a}$ and $\mathsf{val}(\mathsf{tid}, i') \neq \bot$.

Obviously, $i' \notin \mathsf{committed}$ holds, as each fetched instruction is committed only once, and (tid, i') is committed after the load takes place, see (1). All in all, all requirements for the early read from (tid, i') are met, therefore, an early read transition from state s_Z is possible. As shown above, a load from memory transition from the same state leads to $\sigma \notin \mathsf{C}_{\mathsf{power}}(\mathcal{P})$, therefore, (tid, i) reads from the store (tid, i') via an early read transition. □

Lemma 5. *Consider a computation* $\sigma \in \mathsf{C}_{\mathsf{power}}(\mathcal{P})$. *Then a load* (tid, i) *reads a value from a store* (tid', i') *via a load from memory (POW-LOAD) transition iff (1)* $\sigma = \sigma_1 \cdot (\mathsf{prop}, \mathsf{tid}, \mathsf{tid}', i', \mathsf{a}) \cdot \sigma_2 \cdot (\mathsf{load}, \mathsf{tid}, i, \mathsf{a}) \cdot \sigma_3$, *(2)* σ_2 *does not contain events matching* $(\mathsf{prop}, \mathsf{tid}, *, *, \mathsf{a})$, *and (3)* σ_3 *does not contains events matching* $(\mathsf{commit}, \mathsf{tid}, [1..i-1], *, \mathsf{a})$.

Proof. From left to right. Assume the load (tid, i) reads the store (tid', i') via a load from memory transition. Then, the load has read from the latest store to address a propagated to thread tid, i.e., (1) and (2) hold. Assume (3) does not hold — σ_3 contains a commit $(\mathsf{commit}, \mathsf{tid}, i'', *, \mathsf{a})$ and $i'' < i$. Then, (tid, i) reads from the store (tid', i'), which is coherence-order-before the store (tid, i''), which is program-order-before (tid, i). This situation is forbidden by FIN-LD-ST.

From right to left. By (1), (3), and Lemma 4, the load event was not generated by an early read transition. Therefore, the event was generated by a load from memory transition, and the load has taken the value from the latest propagated store to address a, which is, by (1) and (2), (tid', i'). □

3 Robustness

Intuitively, a *trace* $T(\sigma)$ abstracts a program computation σ to the dataflow and control-flow relations between instructions. Formally, the trace of σ is a directed graph $T(\sigma) := (V, \rightarrow_{po}, \rightarrow_{co}, \rightarrow_{src}, \rightarrow_{cf})$ with nodes V and four kinds of arcs. The nodes are instructions together with their thread identifiers and serial numbers (in order to distinguish instructions executed in different threads and the same instruction executed multiple times in the same thread): $V \subseteq (\{\mathsf{init}_\mathsf{a} \mid \mathsf{a} \in \mathsf{ADDR}\} \cup \bigcup_{\mathsf{tid} \in \mathsf{TID}} \{\mathsf{tid}\}) \times \mathbb{N} \times \mathcal{I}_\mathsf{tid}$. The *program order* $\rightarrow_{po}$ is the order in which instructions were fetched in each thread. The *coherence order* $\rightarrow_{co}$ gives the global ordering of writes to each address. The *source order* $\rightarrow_{src}$ shows the store from which a load took its value. The *conflict order* $\rightarrow_{cf}$ shows, for a load, the stores to the same address following the store the load took its value from. We define the *happens-before* relation as $\rightarrow_{hb} := \rightarrow_{po} \cup \rightarrow_{co} \cup \rightarrow_{src} \cup \rightarrow_{cf}$.

Formally, consider a computation $\sigma \in \mathsf{C}_\mathsf{power}(\mathcal{P})$. Let $s_{0Z} \xrightarrow{\sigma} s_Z$ with $s_Z = (\mathsf{ts}, (\mathsf{co}, \mathsf{prop}))$. By Lemma 1, s_Z is uniquely determined. The trace $T(\sigma) := (V, \rightarrow_{po}, \rightarrow_{co}, \rightarrow_{src}, \rightarrow_{cf})$ is defined as follows. Assuming $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{ts}(\mathsf{tid}) = (\mathsf{fetched}, \mathsf{committed}, \mathsf{loaded})$, $i \in [1..|\mathsf{fetched}|]$, and similarly for tid', we have:

$$
\begin{aligned}
V :=& \{(\mathsf{tid}, i, \mathsf{fetched}[i]) \mid \mathsf{tid} \in \mathsf{TID},\ i \in \mathbb{N}\},\\
\rightarrow_{po} :=& \{((\mathsf{tid}, i, \mathsf{fetched}[i]), (\mathsf{tid}, i+1, \mathsf{fetched}[i+1])) \mid\\
&\qquad i \in [1..|\mathsf{fetched}| - 1]\},\\
\rightarrow_{co} :=& \{((\mathsf{tid}, i, \mathsf{fetched}[i]), (\mathsf{tid}', i', \mathsf{fetched}[i'])) \mid\\
&\qquad \mathsf{addr}(\mathsf{tid}, i) = \mathsf{addr}(\mathsf{tid}', i') \text{ and } \mathsf{co}(\mathsf{tid}, i) < \mathsf{co}(\mathsf{tid}', i')\} \cup\\
&\quad \{(\mathsf{init}_a, (\mathsf{tid}', i', \mathsf{fetched}[i'])) \mid a = \mathsf{addr}(\mathsf{tid}', i')\},\\
\rightarrow_{src} :=& \{((\mathsf{tid}, i, \mathsf{fetched}[i]), (\mathsf{tid}', i', \mathsf{fetched}'[i'])) \mid\\
&\qquad (\mathsf{tid}, i) = \mathsf{loaded}'[i']\} \cup\\
&\quad \{(\mathsf{init}_a, (\mathsf{tid}', i', \mathsf{fetched}'[i'])) \mid \mathsf{init}_a = \mathsf{loaded}'(i')\},\\
\rightarrow_{cf} :=& \{(a, b) \mid \exists c\colon c \rightarrow_{src} a \text{ and } c \rightarrow_{co} b\}.
\end{aligned}
$$

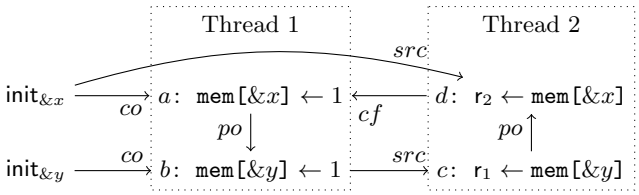

Fig. 2. Trace of computation σ_{MP} from Example 1.

We will also need address $\rightarrow_{addr}$ and data $\rightarrow_{data}$ dependence relations (defined as expected based on $\mathsf{addrdep}$ and $\mathsf{datadep}$).

$$\rightarrow_{addr} := \{((\mathsf{tid}, i, \mathsf{fetched}[i]), (\mathsf{tid}, i', \mathsf{fetched}'[i])) \mid i \in \mathsf{addrdep}(\mathsf{tid}, i')\},$$
$$\rightarrow_{data} := \{((\mathsf{tid}, i, \mathsf{fetched}[i]), (\mathsf{tid}, i', \mathsf{fetched}'[i])) \mid i \in \mathsf{datadep}(\mathsf{tid}, i')\}.$$

Since $\rightarrow_{po}$ includes all the information from the $\mathsf{fetched}$ component of a thread state, $\rightarrow_{addr}$ and $\rightarrow_{data}$ can be reconstructed from $\rightarrow_{po}$ by inspecting the instructions labeling a node. They are therefore not included in the trace explicitly.

The *robustness problem* is, given a program $\mathcal{P}$, to check whether the set of all traces under Power is a subset of all traces under SC: $T_{\mathsf{power}}(\mathcal{P}) \subseteq T_{\mathsf{sc}}(\mathcal{P})$, where $T_{\mathsf{mm}}(\mathcal{P}) := \{T(\sigma) \mid \sigma \in \mathsf{C}_{\mathsf{mm}}(\mathcal{P})\}$ for $\mathsf{mm} \in \{\mathsf{power}, \mathsf{sc}\}$.

Shasha and Snir have shown that a trace belongs to an SC computation iff its happens-before relation is acyclic:

Lemma 6 ([15]). *A program $\mathcal{P}$ is robust against Power iff there is no trace $T \in T_{\mathsf{power}}(\mathcal{P})$ with cyclic $\rightarrow_{hb}$.*

Example 2. The trace of computation σ_{MP} (Figure 2) has a cyclic happens-before relation. By Lemma 6, this means that the program is not robust. Indeed, in no SC computation load d can read 0 whereas c has read 1.

4 Normal-Form Computations

We say that a computation $\tau \in \mathsf{C}_{\mathsf{power}}(\mathcal{P})$ is *in normal form of degree n* if there is a partitioning $\tau = \tau_1 \cdots \tau_n$, such that all fetch events are in τ_1 (NF-A) and events related to different instructions occur in different parts of the computation in the same order (NF-B):

NF-A $(\tau_2 \cdots \tau_n)\downarrow_{\mathsf{fetch}} = \varepsilon$.

NF-B Formally, for $j \in \{1, 2\}$ let e_j, e'_j be events related to instruction (tid_j, i_j). If $\mathsf{e}_1, \mathsf{e}_2 \in \tau_s$ and $\mathsf{e}'_1, \mathsf{e}'_2 \in \tau_{s'}$, then $\mathsf{e}_1 <_{\tau_s} \mathsf{e}_2$ iff $\mathsf{e}'_1 <_{\tau_{s'}} \mathsf{e}'_2$.

In the rest of this section we prove the following theorem:

Theorem 1. *A program is robust iff it has no normal-form computations of degree $|\mathcal{P}| + 3$ with cyclic happens-before relation.*

Consider a computation $\sigma \in \mathsf{C_{mm}}(\mathcal{P})$. By $\sigma \setminus (\mathsf{tid}, i)$ we denote the computation obtained from σ by deleting all events related to the i'th fetched instruction in thread tid.

Lemma 7. *Consider a non-empty computation $\sigma \in \mathsf{C}_{power}(\mathcal{P})$. Then there is a (tid_x, i_x), such that $\sigma' = \sigma \setminus (\mathsf{tid}_x, i_x)$ satisfies $|\sigma'| < |\sigma|$ and $\sigma' \in \mathsf{C}_{power}(\mathcal{P})$.*

Proof. Consider the last fetched instruction in each thread. If among such instructions there is a non-store instruction, delete it: its result cannot be used by any other instruction. If all these instructions are stores, delete the one, on which (1) no load or store depends via $(\rightarrow_{src} \cup \rightarrow_{data})^+ \cdot \rightarrow_{addr}$, and (2) no condition depends via $(\rightarrow_{src} \cup \rightarrow_{data})^+$.

Towards a contradiction, assume there is no such store. Consider a last fetched (store) instruction in a thread tid_1: (tid_1, i_1). Case 1: there is a load or a store (tid_2, i_2') whose address depends on (tid_1, i_1). Case 2: there is a condition (tid_2, i_2') whose value depends on (tid_1, i_1). Consider the last fetched instruction in thread tid_2: (tid_2, i_2). It must be a store, and it must have been committed after (tid_1, i_1): a store can only be committed after all loads and stores fetched before it have their addresses determined (Case 1) and after all preceding conditions are committed (Case 2).

Continuing the reasoning, for any last fetched instruction in a thread (tid_j, i_j) there is a last instruction in a different thread $(\mathsf{tid}_{j+1}, i_{j+1})$ which must have been committed later. Taking into account finiteness of the number of threads, we get a contradiction. □

Fix a program $\mathcal{P}$. Consider a shortest Power computation $\alpha \in \mathsf{C_{power}}(\mathcal{P})$ with cyclic $\rightarrow_{hb}$. Let (tid_x, i_x) be the instruction determined by Lemma 7. Let $\alpha := \alpha_1 \cdot \mathsf{x}_1 \cdot \alpha_2 \cdot \mathsf{x}_2 \cdots \alpha_n$, where $\{\mathsf{x}_1 \ldots \mathsf{x}_{n-1}\}$ are the events related to the i_x'th instruction fetched in thread tid_x. Then $\alpha \setminus (\mathsf{tid}_x, i_x) := \alpha' := \alpha_1 \cdot \alpha_2 \cdots \alpha_n$. Since α' is shorter than α, its $\rightarrow_{hb}$ is acyclic. Therefore, there is a computation $\beta \in \mathsf{C_{sc}}(\mathcal{P})$ with $T(\beta) = T(\alpha')$.

Computations β and α' consist of the same fetch, load, and commit events: fetch events are determined by $\rightarrow_{po}$; address component a of load and store commit events is determined by $\rightarrow_{addr}$, $\rightarrow_{data}$ (derivable from $\rightarrow_{po}$), and $\rightarrow_{src}$; since $\rightarrow_{co}$ is the same for both computations, we can assume that matching store commit events have the same value of coherence key k. Notably, β can have more propagate events than α' as Power semantics does not guarantee that all stores are propagated to all threads. Now we reorder events in each part α_j of α in the way they follow in β. This gives a computation $\gamma := \beta{\downarrow}\alpha_1 \cdot \mathsf{x}_1 \cdot \beta{\downarrow}\alpha_2 \cdot \mathsf{x}_2 \cdots \beta{\downarrow}\alpha_n$. In the rest of the section we show that γ is a valid Power computation of program $\mathcal{P}$ and has the same trace as α.

Lemma 8. *For all $\mathsf{tid} \in \mathsf{TID}$ holds $\alpha{\downarrow}\mathsf{fetch}{\downarrow}\mathsf{tid} = \gamma{\downarrow}\mathsf{fetch}{\downarrow}\mathsf{tid}$.*

Proof. Since $T(\beta) = T(\alpha')$, by definition of α and properties of projection, for any $\mathsf{tid} \in \mathsf{TID}$ we have

$$\begin{aligned}
\alpha{\downarrow}\mathsf{fetch}{\downarrow}\mathsf{tid} &= \alpha_1{\downarrow}\mathsf{fetch}{\downarrow}\mathsf{tid} \cdot \mathsf{x}_1{\downarrow}\mathsf{fetch}{\downarrow}\mathsf{tid} \cdots \alpha_n{\downarrow}\mathsf{fetch}{\downarrow}\mathsf{tid} \\
&= \cdots (\beta{\downarrow}\mathsf{fetch}{\downarrow}\mathsf{tid}){\downarrow}(\alpha_i{\downarrow}\mathsf{fetch}{\downarrow}\mathsf{tid}) \cdot \mathsf{x}_i{\downarrow}\mathsf{fetch}{\downarrow}\mathsf{tid} \cdots \\
&= \beta{\downarrow}\alpha_1{\downarrow}\mathsf{fetch}{\downarrow}\mathsf{tid} \cdot \mathsf{x}_1{\downarrow}\mathsf{fetch}{\downarrow}\mathsf{tid} \cdots \beta{\downarrow}\alpha_n{\downarrow}\mathsf{fetch}{\downarrow}\mathsf{tid} \\
&= (\beta{\downarrow}\alpha_1 \cdot \mathsf{x}_1 \cdots \beta{\downarrow}\alpha_n){\downarrow}\mathsf{fetch}{\downarrow}\mathsf{tid} \\
&= \gamma{\downarrow}\mathsf{fetch}{\downarrow}\mathsf{tid}.
\end{aligned}$$

□

Lemma 9. *Consider some* (tid, i) *and* (tid', i')*. Let* $P(\sigma) := \mathsf{true}$ *if requirements (1)–(2) from Lemma 4 or (1)–(3) from Lemma 5 hold for* σ*, and* $P(\sigma) := \mathsf{false}$ *otherwise. Then, if* $P(\alpha)$ *then* $P(\gamma)$*.*

Proof. The proof is a case consideration: which of the two cases holds in the definition of P hold for σ, for α, and whether the distinguished load and commit events are located in the same part α_j. We consider two of the cases. The other are similar.

Assume requirements (1)–(2) from Lemma 4 hold for α and requirements (1)–(3) from Lemma 5 holds for sequentially consistent computation β. If load and commit events are in the same part, then $\alpha = \alpha_1 \cdot \mathsf{x}_1 \cdots (\alpha'_j \cdot b \cdot \alpha''_j \cdot c \cdot d \cdot \alpha'''_j) \cdot \mathsf{x}_j \cdots \alpha_n$, $\beta = \beta_1 \cdots c \cdot d \cdot \beta_2 \cdot b \cdot \beta_3$, where $b = (\mathsf{load}, \mathsf{tid}, i, \mathsf{a})$, $c = (\mathsf{commit}, \mathsf{tid}, i')$, $d = (\mathsf{prop}, \mathsf{tid}, \mathsf{tid}, i', \mathsf{a})$, $i' < i$. Consequently, $\gamma = \beta{\downarrow}\alpha_1 \cdot \mathsf{x}_1 \cdots \beta{\downarrow}\alpha_j \cdot \mathsf{x}_j \cdots \beta{\downarrow}\alpha_n = \beta{\downarrow}\alpha_1 \cdot \mathsf{x}_1 \cdots (\beta_1{\downarrow}\alpha_j \cdot d \cdot \beta_2{\downarrow}\alpha_j \cdot b \cdot \beta_3{\downarrow}\alpha_j) \cdot \mathsf{x}_j \cdots \beta{\downarrow}\alpha_n$ — looks like a read from memory situation. We check requirements (1)–(3) of Lemma 5 then. First, $\beta_2{\downarrow}\alpha_j$ must have no prop events to thread tid with the address a — holds as β_2 does not have them. Second, $\beta_3{\downarrow}\alpha_j$ must have no commits of earlier writes in thread tid — holds as β_3 does not have them. Third, $\beta{\downarrow}\alpha_l = (\beta_1 \cdot \beta_2 \cdot \beta_3){\downarrow}\alpha_l$, $l \in [i+1..n]$ must have no commit events for stores with indices $[1..i-1]$, the same address and thread id. Consider $\beta_1{\downarrow}\alpha_l$ — if it has such an event e, then two stores to the same address, e and c, are committed in different order in α' and β, which is impossible due to $T(\alpha') = T(\beta)$. Consider $\beta_2{\downarrow}\alpha_l$ — it does not have such an event, because β_2 does not have prop events to address a, therefore, it does not have commits of own stores there too. Consider $\beta_3{\downarrow}\alpha_l$ — it does not have such an event, because β_3 does not. Finally, none of x_l events, $l \in [i+1..n-1]$, must be a commit of earlier writes in thread tid — holds, as these events belong to the last fetched instruction of a thread.

Consider the case when load and commit events are in different parts, i.e. $\alpha = \alpha_1 \cdot \mathsf{x}_1 \cdots (\alpha'_j \cdot b \cdot \alpha''_j) \cdots (\alpha'_k \cdot c \cdot d \cdot \alpha''_k) \cdots \alpha_n$, $\beta = \beta_1 \cdot c \cdot d \cdot \beta_2 \cdot b \cdot \beta_3$, where b, c, d are defined as before and $i' < i$. Then, $\gamma = \beta \downarrow \alpha_1 \cdot \mathsf{x}_1 \cdots \beta \downarrow \alpha_j \cdot \mathsf{x}_j \cdots \beta \downarrow \alpha_k \cdots \beta \downarrow \alpha_n = \beta{\downarrow}\alpha_1 \cdot \mathsf{x}_1 \cdots \beta_1{\downarrow}\alpha_j \cdot \beta_2{\downarrow}\alpha_j \cdot b \cdot \beta_3{\downarrow}\alpha_j \cdot \mathsf{x}_j \cdots \beta_1{\downarrow}\alpha_k \cdot c \cdot d \cdot \beta_2{\downarrow}\alpha_k \cdot \beta_3{\downarrow}\alpha_k \cdot \mathsf{x}_k \cdots \beta{\downarrow}\alpha_n$ — looks like an early read case. Therefore, one must check that $\beta_2{\downarrow}\alpha_k \cdot \beta_3{\downarrow}\alpha_k \cdot \mathsf{x}_k \cdots \beta{\downarrow}\alpha_n$ has no commit events matching $(\mathsf{commit}, \mathsf{tid}, [i'+1..i-1], *, \mathsf{a})$. Consider $\beta_2 \downarrow \alpha_k$ — does not have such events, because they would be immediately followed by a prop event to thread tid and

address a, which contradicts requirement (2) of Lemma 5. Consider $\beta_3 \downarrow \alpha_k$ — does not have such events, because β_3 does not have them by requirement (3) of Lemma 5. Consider $\beta \downarrow \alpha_l$, $l \in [j+1..n]$ — does not have such events, because α_l do not have them by requirement (2) of Lemma 4. Finally, x_l, $l \in [j+1..n-1]$ belong to the last fetched instruction of a thread, therefore do not contain the described commit events. □

Lemma 10. $\gamma \in \mathsf{C}_{\mathit{power}}(\mathcal{P})$.

Proof. We proceed by induction. Assume (1) $\gamma = \gamma_1 \cdot \mathsf{e} \cdot \gamma_2$, (2) $s_{0Z} \xrightarrow{\gamma_1} s_Z$, and (3) all loads satisfied in γ_1 have read from the same stores as in α. We show that $s_{0Z} \xrightarrow{\gamma_1 \cdot \mathsf{e}} s_Z{}'$ and all loads satisfied in $\gamma_1 \cdot \mathsf{e}$ have read from the same stores as in α. Let $s_Z = (\mathsf{ts}, s_Y)$ and $\mathsf{ts}(\mathsf{tid}) = (\mathsf{fetched}, \mathsf{committed}, \mathsf{loaded})$. Consider the event e.

($\mathbf{fetch}, \mathbf{tid}, i$) A transition labeled by e from state s_Z is feasible due to Lemma 8 and the fact that feasibility of a fetch transition is conditioned solely on the previous fetch transition with the same thread id.

($\mathbf{load}, \mathbf{tid}, i, \mathbf{a}$) For the transition to be feasible, $\mathsf{addr}(i) = \mathsf{a}$ must hold. In order to have $\mathsf{addr}(\mathsf{tid}, i) \neq \bot$, all loads in thread tid, on which $\mathsf{addr}(\mathsf{tid}, i)$ depends, must be satisfied. Note that these loads are the same in α and γ due to Lemma 8. Since $\alpha \in \mathsf{C}_{\mathsf{power}}(\mathcal{P})$, these load events occurred before e in α. Let e' be one of these load events. If $\mathsf{e}' \in \alpha_i$ and $\mathsf{e} \in \alpha_j$, $i < j$, or $\mathsf{e}' \in \{\mathsf{x}_i \mid i \in [1..n-1]\}$, or $\mathsf{e} \in \{\mathsf{x}_i \mid i \in [1..n-1]\}$, then e' and e are located in γ in the same order. If $\mathsf{e}', \mathsf{e} \in \alpha_i$, then $\mathsf{e}', \mathsf{e} \in \beta$. Since the $\rightarrow_{po}$ components of $T(\alpha)$ and $T(\beta)$ match up to a single deleted arc, e' and e are located in β (therefore, in $\beta \downarrow \alpha_i$ and γ) in this order. By inductive assumption (3) and the fact that functions in FUN are deterministic, $\mathsf{addr}(\mathsf{tid}, i) = \mathsf{a}$ holds.
Assume the load (tid, i) has read from a store (tid', i') in α. Then, by Lemmas 4, 5, 9, either conditions (1)–(3) of Lemma 5 hold, or conditions (1)–(2) of Lemma 4 hold. In the former case, $(\mathsf{prop}, \mathsf{tid}, \mathsf{tid}', i', \mathsf{a})$ is the last prop event to tid with address a, therefore, a load from memory transition reading (tid', i') is feasible from state s_Z. In the latter case, (tid', i') is the latest non-committed store to address a, and an early read transition reading (tid', i') is possible. The proof that $\mathsf{addr}(\mathsf{tid}', i') \neq \bot$ is similar to the proof that $\mathsf{addr}(\mathsf{tid}, i) \neq \bot$.

($\mathbf{commit}, \mathbf{tid}, i$) The proof of $\mathsf{addr}(\mathsf{tid}, i) \neq \bot$ and $\mathsf{val}(\mathsf{tid}, i) \neq \bot$ is similar to the proof of $\mathsf{addr}(\mathsf{tid}, i) \neq \bot$ in the previous case. If $\mathsf{fetched}[i]$ is a load or a store, there must be no preceding loads and stores to unknown addresses, which holds and can be proven in a similar way. If $\mathsf{fetched}[i]$ is a load, requirement $\mathsf{loaded}[i] \neq \bot$ holds for the same reasons. If $\mathsf{fetched}[i]$ is a conditional, requirement $\mathsf{val}(\mathsf{tid}, i) \neq 0$ holds by inductive assumption (3), the fact that functions in FUN are deterministic, and the fact that $\alpha \in \mathsf{C}_{\mathsf{power}}(\mathcal{P})$.

($\mathbf{commit}, \mathbf{tid}, i, \mathbf{k}, \mathbf{a}$) Value k is unique, since it was unique in α, and α and γ consist of the same commit events. We check $\mathsf{co}(\mathsf{prop}(\mathsf{tid}, \mathsf{a})) < \mathsf{k}$. Assume it does not hold. Then, there is $\mathsf{e}' = (\mathsf{prop}, \mathsf{tid}, \mathsf{tid}', i', \mathsf{a})$, where $\mathsf{co}(\mathsf{tid}', i') > \mathsf{k}$,

and e', e are located in γ in this order. If $\mathsf{e}' \in \alpha_i$, $\mathsf{e} \in \alpha_j$, $i < j$, or $\mathsf{e}' \in \{\mathsf{x}_i \mid i \in [1..n-1]\}$, or $\mathsf{e} \in \{\mathsf{x}_i \mid i \in [1..n-1]\}$, these events are located in α in this order, which contradicts $\alpha \in \mathsf{C}_{\mathsf{power}}(\mathcal{P})$. If $\mathsf{e}', \mathsf{e} \in \alpha_i$, these events are located in β in this order, which contradicts $\beta \in \mathsf{C}_{\mathsf{power}}(\mathcal{P})$.
This transition is immediately followed by a prop transition in γ, since it did so in α and β (unless $e \in \{\mathsf{x}_i \mid i \in [1..n-1]\}$, which is a simpler case), and by properties of projection.

($\mathbf{prop}, \mathbf{tid}, \mathbf{tid}', i', \mathbf{a}$) The requirement $\mathsf{co}(\mathsf{prop}(\mathsf{tid}, \mathsf{a})) < \mathsf{co}(\mathsf{tid}', i')$ is proven similarly to $\mathsf{co}(\mathsf{prop}(\mathsf{tid}, \mathsf{a})) < \mathsf{k}$ in the previous case.

As shown above, $s_{0Z} \xrightarrow{\gamma} s_Z$. What is left to check, is that $s_Z \in F_Z$. The requirement that all fetched instructions are committed trivially holds: β includes the same commit events as α', therefore, by definition, γ contains the same commit events as α. The other two requirements that loads and stores agree with the coherence order hold due to Lemma 8, the inductive assumption (3), and the fact that α and γ consist of the same commit events (i.e. the coherence keys of matching stores are equal in these computations). □

Lemma 11. $T(\gamma) = T(\alpha)$

Proof. Equality of $\rightarrow_{po}$ follows from Lemma 8. Equality of source relation follows from Lemmas 4, 5, 9, 10. Store order is determined by a and k components of store commit events. Since computations α and γ consist of the same commit events, the $\rightarrow_{co}$ relations in the traces of α and γ are the same. □

Lemma 12. $\gamma \in \mathsf{C}_{\mathsf{power}}(\mathcal{P})$ *and* $T(\gamma) = T(\alpha)$.

Proof. Corollary of Lemmas 10 and 11. □

Without loss of generality we may assume that all fetch events of α are located within $\alpha_1 \cdot \mathsf{x}_1$: every thread can always first fetch all instructions and in the rest of the computation only execute them; such a reordering does not change the trace. Also, note that the maximal number of events an instruction can generate is $|\mathcal{P}| + 2$. This bound is achieved by a store that is fetched, committed, and propagated to all threads. Then the following lemma holds:

Lemma 13. *Computation γ is in normal form of degree $|\mathcal{P}| + 3$.*

Proof. By definition of γ and properties of projection. □

Together with Lemma 6 this proves Theorem 1.

Example 3. Consider $\alpha := \mathsf{fetch}(c) \cdot \mathsf{fetch}(d) \cdot \mathsf{fetch}(a) \cdot \cancel{\mathsf{fetch}(b)} \cdot \mathsf{commit}(a) \cdot \mathsf{prop}(a,1) \cdot \cancel{\mathsf{commit}(b)} \cdot \cancel{\mathsf{prop}(b,1)} \cdot \cancel{\mathsf{prop}(b,2)} \cdot \mathsf{load}(c) \cdot \mathsf{load}(d) \cdot \mathsf{commit}(d) \cdot \mathsf{commit}(c)$, which is essentially σ_{MP} with fetch events moved to the front. We cancel the x_i events (crossed out) related to the store instruction b, as b is the last instruction of thread 1 and no address depends on it (we could also cancel the events of d instead). Therefore, $\alpha_1 := \mathsf{fetch}(c) \cdot \mathsf{fetch}(d) \cdot \mathsf{fetch}(a)$, $\alpha_2 := \mathsf{commit}(a) \cdot \mathsf{prop}(a,1)$, $\alpha_3 := \alpha_4 := \varepsilon$, $\alpha_5 := \mathsf{load}(c) \cdot \mathsf{load}(d) \cdot \mathsf{commit}(d) \cdot \mathsf{commit}(c)$, and $\alpha' :=$

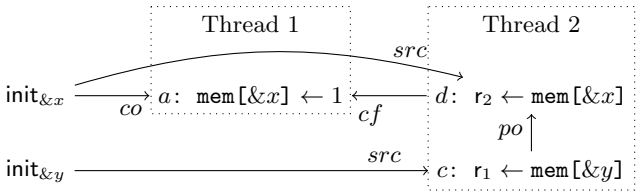

Fig. 3. Trace of the computations α' and β from Example 3.

$\alpha_1 \cdot \alpha_2 \cdot \alpha_3 \cdot \alpha_4 \cdot \alpha_5$. The trace of α' is shown in Figure 3. The SC computation with the same trace is $\beta := \mathsf{fetch}(c) \cdot \mathsf{load}(c) \cdot \mathsf{commit}(c) \cdot \mathsf{fetch}(d) \cdot \mathsf{load}(d) \cdot \mathsf{commit}(d) \cdot \mathsf{fetch}(a) \cdot \mathsf{commit}(a) \cdot \mathsf{prop}(a,1) \cdot \mathsf{prop}(a,2)$. The normal-form computation is $\gamma := \beta{\downarrow}\alpha_1 \cdot \mathsf{x}_1 \cdots \beta{\downarrow}\alpha_5 = (\mathsf{fetch}(c) \cdot \mathsf{fetch}(d) \cdot \mathsf{fetch}(a)) \cdot \mathsf{fetch}(b) \cdot (\mathsf{commit}(a) \cdot \mathsf{prop}(a,1)) \cdot \mathsf{commit}(b) \cdot \mathsf{prop}(b,1) \cdot \mathsf{prop}(b,2) \cdot (\mathsf{load}(c) \cdot \mathsf{commit}(c) \cdot \mathsf{load}(d) \cdot \mathsf{commit}(d))$. It is feasible and has the same trace as α and σ_{MP} (Figure 2).

5 From Normal-Form Computations to Emptiness

We now reduce robustness to language emptiness. First, we define a multiheaded automaton capable of generating all normal-form computations of a program. Next, we intersect it with regular languages that check cyclicity of the happens-before relation. Altogether, the program is robust iff the intersection is empty.

5.1 Generating Normal-Form Computations

To generate all normal-form computations, we use so-called multiheaded automata [8]. Essentially, a multiheaded automaton generates a computation $\sigma_1 \dots \sigma_n$ by simultaneously generating its parts σ_i. The automaton has a head for each part, and the transition relation defines the head producing an event. Formally, an *n-headed automaton over Σ* is an automaton operating on the extended alphabet $[1..n] \times \Sigma$: $A = (S, [1..n] \times \Sigma, \Delta, s_0, F)$. The *language of A* is $\mathcal{L}(A) := \{\mathsf{second}(\sigma{\downarrow}(\{1\}\times\Sigma) \cdots \sigma{\downarrow}(\{n\}\times\Sigma)) \mid s_0 \xrightarrow{\sigma} s \text{ for some } s \in F\}$, where $\mathsf{second}((a_1,b_1)\cdots(a_m,b_m)) := b_1 \cdots b_m$. Multiheaded automata are closed under intersection with regular languages. Moreover, for finite multiheaded automata language emptiness is NL-complete [8]:

Lemma 14 ([8]). *Consider an n-headed automaton U and a finite automaton V over a common alphabet Σ. There is an n-headed automaton W with $\mathcal{L}(W) = \mathcal{L}(U) \cap \mathcal{L}(V)$ with the number of states $|S_W| \leq |S_U| \cdot |S_V|^{2n} + 1$.*

Lemma 15 ([8]). *Emptiness for n-headed automata is* NL*-complete.*

We will generate all normal-form computations of program $\mathcal{P}$ with the n-headed automaton $M(\mathcal{P}) := (S_M, \mathsf{E}, \Delta_M, s_{0M}, F_M)$, where $n := |\mathcal{P}| + 3$. The automaton generates all events related to a single instruction in one shot, but,

possibly, in different parts of the computation. All fetch events are generated in the first part of the computation. In order to generate them, the automaton keeps track of the destination state of the last fetched instruction in each thread (component ctrl-state of the automaton state).

Each instruction can only read the last value written to a register. Therefore, the automaton only needs to remember |REG| register values per thread (component reg-value). However, an instruction cannot be executed until the values of all registers that it reads become known. To obey this restriction, the automaton memorizes the part of the computation in which the register value gets computed (reg-comp-head). For example, while handling an assignment $\mathsf{r}_1 \leftarrow \mathsf{r}_1 + \mathsf{r}_2$, the automaton learns that the new value of r_1 is the sum of the current values of r_1 and r_2. It also remembers that this value is available no earlier than the current values of r_1 and r_2 are computed. Similarly, the automaton remembers the parts of the computation in which the addresses of load and store instructions become known (addr-comp-head), and certain kinds of instructions get committed (reg-comm-head, assume-comm-head, addr-comm-head).

The automaton has to keep a separate memory state for each thread and for each part of the computation. The memory state of a thread in a part is updated when a store instruction gets propagated to this thread in this part. When a load instruction is handled, the automaton chooses a part where the load event takes place and uses the memory state of that part. Besides the memory valuation (mem-value), the memory state includes coherence keys (last-key) to guarantee that the generated computation respects the coherence order.

When starting the computation, the automaton non-deterministically guesses the memory valuations and coherence keys for all parts of the computation (except the first one). Upon termination, the automaton checks that the parts of the computation generated by each head fit together at the concatenation points. This ensures the overall computation is valid for the program. The trick is to remember the guess of the initial memory valuations and coherence keys in immutable components of the automaton state ($\mathsf{mem\text{-}value}_g$, $\mathsf{last\text{-}key}_g$). The final states require that the current memory state in part h of the computation coincides with the guessed initial state in part $\mathsf{h}+1$.

State space A state from S_M (except the special initial state s_{0M}) includes the following information:

- ctrl-state(tid) gives the current control state of thread tid.
- reg-comp-head(tid, r) gives the part in which last value assigned to register r in thread tid gets computed.
- reg-value(tid, r) gives this computed value.
- reg-comm-head(tid, r) gives the part in which the last instruction assigning a value to register r in thread tid gets committed.
- assume-comm-head(tid) gives the part in which the latest fetched condition in thread tid is committed.
- mem-value(tid, a, h) gives the value of the last write to a propagated to thread tid in the part h or earlier.

– $\mathsf{last\text{-}key}(\mathsf{tid}, \mathsf{a}, \mathsf{h})$ gives the coherence key of the last write to a propagated to thread tid in the part h or earlier.
– $\mathsf{mem\text{-}value}_g$, $\mathsf{last\text{-}key}_g$ are immutable copies of the guessed values of the previous two components (see MH-GUESS below).
– $\mathsf{early\text{-}mem\text{-}value}(\mathsf{tid}, \mathsf{a}, \mathsf{h})$ gives the value written by the last fetched store to a which is still in-flight in the part h of computation, $\bot$ if there is no such store, $\top$ if the value of the store is unknown or there is a later in-flight store in this part with an unknown address.
– $\mathsf{addr\text{-}comp\text{-}head}(\mathsf{tid})$ gives the leftmost part of the computation, in which the addresses of all already fetched memory accesses are computed.
– $\mathsf{addr\text{-}comm\text{-}head}(\mathsf{tid}, \mathsf{a})$ gives the rightmost part of the computation having a commit to address a by thread tid.
– $\mathsf{instr\text{-}count}(\mathsf{tid})$ gives the number of instructions fetched in thread tid.

The initial state s_{0M} does not contain any information.

Transition relation We define transitions by specifying the new (primed) values of the state components and the label λ of the transition. First, we define the transition guessing the initial memory state in each part of the computation:

MH-GUESS Assume the current state is s_{0M}. Then, there are transitions to the states satisfying $\mathsf{ctrl\text{-}state}' := \lambda\mathsf{tid}.q_{0\mathsf{tid}}$, $\mathsf{reg\text{-}comp\text{-}head}' := \lambda\mathsf{tid}.\lambda\mathsf{r}.1$, $\mathsf{reg\text{-}value}' := \lambda\mathsf{tid}.\lambda\mathsf{r}.0$, $\mathsf{reg\text{-}comm\text{-}head}' := \lambda\mathsf{tid}.\lambda\mathsf{r}.1$, $\mathsf{assume\text{-}comm\text{-}head}' := \lambda\mathsf{tid}.1$, $\mathsf{early\text{-}mem\text{-}value}' := \lambda\mathsf{tid}.\lambda\mathsf{a}.\lambda\mathsf{h}.\bot$, $\mathsf{mem\text{-}value}' = \mathsf{mem\text{-}value}'_g$, $\mathsf{last\text{-}key}' = \mathsf{last\text{-}key}'_g$, $\mathsf{addr\text{-}comp\text{-}head}' := \lambda\mathsf{tid}.1$, $\mathsf{addr\text{-}comm\text{-}head}' := \lambda\mathsf{tid}.\lambda\mathsf{a}.1$, $\mathsf{instr\text{-}count}' := \lambda\mathsf{tid}.0$. Also, $\mathsf{mem\text{-}value}'(\mathsf{tid}, \mathsf{a}, 1) := 0$, $\mathsf{last\text{-}key}'(\mathsf{tid}, \mathsf{a}, 1) := 0$ for all $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{a} \in \mathsf{ADDR}$. Moreover, $\mathsf{last\text{-}key}'(\mathsf{tid}, \mathsf{a}, \mathsf{h}) \leq \mathsf{last\text{-}key}'(\mathsf{tid}, \mathsf{a}, \mathsf{h}+1)$ for $\mathsf{h} \in [1..n-1]$, $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{a} \in \mathsf{ADDR}$ (we assume $\mathsf{last\text{-}key}'(\mathsf{tid}, \mathsf{a}, n) := \infty$). $\lambda := \varepsilon$.

Fix a state s_M. We overload $\mathsf{eval}(\mathsf{tid}, e)$ to mean the value of expression e for the valuation of registers defined by $\lambda\mathsf{r}.\mathsf{reg\text{-}value}(\mathsf{tid}, \mathsf{r})$.

Let $\mathsf{HEAD} := [1..n]$. Let $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{ctrl\text{-}state}(\mathsf{tid}) = q_1$, $\mathsf{instr} = q_1 \xrightarrow{\mathsf{cmd}} q_2 \in \mathcal{I}_{\mathsf{tid}}$. Let $\mathsf{h}_1 := 1$. Let $\mathsf{h}_2 \in \mathsf{HEAD}$, $\mathsf{h}_2 \geq \mathsf{h}_1$, $\mathsf{h}_2 \geq \mathsf{reg\text{-}comp\text{-}head}(\mathsf{tid}, \mathsf{r})$ for each register r read in cmd. Let $\mathsf{h}_3 \in \mathsf{HEAD}$, $\mathsf{h}_3 \geq \mathsf{h}_2$, $\mathsf{h}_3 \geq \mathsf{reg\text{-}comm\text{-}head}(\mathsf{tid}, \mathsf{r})$ for each register r read in cmd, $\mathsf{h}_3 \geq \mathsf{assume\text{-}comm\text{-}head}(\mathsf{tid})$. Let $i := \mathsf{instr\text{-}count}(\mathsf{tid}) + 1$ and $\mathsf{instr\text{-}count}' := \mathsf{instr\text{-}count}[\mathsf{tid} \hookleftarrow i]$. Depending on the type of cmd, there are the following transitions from s_M labeled by events λ:

MH-ASSIGN $\mathsf{cmd} = \mathsf{r} \leftarrow e_{\mathsf{v}}$. Let $\mathsf{v} := \mathsf{eval}(\mathsf{tid}, e_{\mathsf{v}})$. Then $\mathsf{reg\text{-}value}' := \mathsf{reg\text{-}value}[(\mathsf{tid}, \mathsf{r}) \hookleftarrow \mathsf{v}]$, $\mathsf{reg\text{-}comp\text{-}head}' := \mathsf{reg\text{-}comp\text{-}head}[(\mathsf{tid}, \mathsf{r}) \hookleftarrow \mathsf{h}_2]$, $\mathsf{reg\text{-}comm\text{-}head}' := \mathsf{reg\text{-}comm\text{-}head}[(\mathsf{tid}, \mathsf{r}) \hookleftarrow \mathsf{h}_3]$. $\lambda := (\mathsf{h}_1, \mathsf{fetch}, \mathsf{tid}, \mathsf{instr}) \cdot (\mathsf{h}_3, \mathsf{commit}, \mathsf{tid}, i)$.

MH-ASSUME $\mathsf{cmd} = \mathtt{assume}(e_{\mathsf{v}})$. Let $\mathsf{eval}(\mathsf{tid}, e_{\mathsf{v}}) \neq 0$. Then $\mathsf{assume\text{-}comm\text{-}head}' := \mathsf{assume\text{-}comm\text{-}head}[\mathsf{tid} \hookleftarrow \mathsf{h}_3]$. $\lambda := (\mathsf{h}_1, \mathsf{fetch}, \mathsf{tid}, \mathsf{instr}) \cdot (\mathsf{h}_3, \mathsf{commit}, \mathsf{tid}, i)$.

MH-LOAD $\mathsf{cmd} = \mathsf{r} \leftarrow \texttt{mem}[e_\mathsf{a}]$. Let $\mathsf{a} := \mathsf{eval}(\mathsf{tid}, e_\mathsf{a})$. Let $\mathsf{h}_3 \geq \mathsf{addr\text{-}comm\text{-}head}(\mathsf{tid}, \mathsf{a})$. If $\mathsf{early\text{-}mem\text{-}value}(\mathsf{tid}, \mathsf{a}) = \bot$, let $\mathsf{v} := \mathsf{mem\text{-}value}(\mathsf{tid}, \mathsf{a}, \mathsf{h}_2)$ (load from memory case). Otherwise, let $\mathsf{v} := \mathsf{early\text{-}mem\text{-}value}(\mathsf{tid}, \mathsf{a}, \mathsf{h}_2)$ and assume $\mathsf{v} \neq \top$ (early read case). Then $\mathsf{reg\text{-}value}' := \mathsf{reg\text{-}value}[(\mathsf{tid}, \mathsf{r}) \hookleftarrow \mathsf{v}]$, $\mathsf{reg\text{-}comp\text{-}head}' := \mathsf{reg\text{-}comp\text{-}head}[(\mathsf{tid}, \mathsf{r}) \hookleftarrow \mathsf{h}_2]$, $\mathsf{reg\text{-}comm\text{-}head}' := \mathsf{reg\text{-}comm\text{-}head}[(\mathsf{tid}, \mathsf{r}) \hookleftarrow \mathsf{h}_3]$, $\mathsf{addr\text{-}comp\text{-}head}' := \mathsf{addr\text{-}comp\text{-}head}[\mathsf{tid} \hookleftarrow \max\{\mathsf{addr\text{-}comp\text{-}head}(\mathsf{tid}), \mathsf{h}_2\}]$, $\mathsf{addr\text{-}comm\text{-}head}' := \mathsf{addr\text{-}comm\text{-}head}[(\mathsf{tid}, \mathsf{a}) \hookleftarrow \mathsf{h}_3]$. $\lambda := (\mathsf{h}_1, \mathsf{fetch}, \mathsf{tid}, \mathsf{instr}) \cdot (\mathsf{h}_2, \mathsf{load}, \mathsf{tid}, i, \mathsf{a}) \cdot (\mathsf{h}_3, \mathsf{commit}, \mathsf{tid}, i)$.

MH-STORE $\mathsf{cmd} = \texttt{mem}[e_\mathsf{a}] \leftarrow e_\mathsf{v}$. Let $\mathsf{a} := \mathsf{eval}(\mathsf{tid}, e_\mathsf{a})$. Assume $\mathsf{h}_3 \geq \mathsf{addr\text{-}comp\text{-}head}(\mathsf{tid})$, $\mathsf{h}_3 \geq \mathsf{addr\text{-}comm\text{-}head}(\mathsf{tid}, \mathsf{a})$. Let $\mathsf{v} := \mathsf{eval}(\mathsf{tid}, e_\mathsf{v})$. Let $\mathsf{k} \in \mathbb{Q}$, $\mathsf{k} \neq \mathsf{last\text{-}key}(\mathsf{tid}, \mathsf{a}, \mathsf{h})$ for any $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{a} \in \mathsf{ADDR}$, $\mathsf{h} \in \mathsf{HEAD}$. Then $\mathsf{early\text{-}mem\text{-}value}' := \mathsf{early\text{-}mem\text{-}value}[(\mathsf{tid}, \mathsf{a}, [\mathsf{h}_1..\mathsf{h}_2 - 1]) \hookleftarrow \top), (\mathsf{tid}, \mathsf{a}, [\mathsf{h}_2..\mathsf{h}_3 - 1]) \hookleftarrow \mathsf{v}]$. We also set $\mathsf{early\text{-}mem\text{-}value}' := \mathsf{early\text{-}mem\text{-}value}'[(\mathsf{tid}, \mathsf{a}', \mathsf{h}) \hookleftarrow \top]$ for all $\mathsf{a}' \in \mathsf{ADDR} \setminus \{\mathsf{a}\}$, $\mathsf{h} \in [\mathsf{h}_1..\mathsf{h}_2 - 1]$ with $\mathsf{early\text{-}mem\text{-}value}(\mathsf{tid}, \mathsf{a}', \mathsf{h}) \in \mathsf{DOM}$. We define $\mathsf{addr\text{-}comp\text{-}head}' := \mathsf{addr\text{-}comp\text{-}head}[\mathsf{tid} \hookleftarrow \max\{\mathsf{addr\text{-}comp\text{-}head}(\mathsf{tid}), \mathsf{h}_2\}]$, $\mathsf{addr\text{-}comm\text{-}head}' := \mathsf{addr\text{-}comm\text{-}head}[(\mathsf{tid}, \mathsf{a}) \hookleftarrow \mathsf{h}_3]$. Let $T \subseteq \mathsf{TID} \setminus \{\mathsf{tid}\}$, initially $\mathsf{mem\text{-}value}' := \mathsf{mem\text{-}value}$, $\mathsf{last\text{-}key}' := \mathsf{last\text{-}key}$, and $\lambda := (\mathsf{h}_1, \mathsf{fetch}, \mathsf{tid}, \mathsf{instr}) \cdot (\mathsf{h}_3, \mathsf{commit}, \mathsf{tid}, i, \mathsf{k}, \mathsf{a})$. For $\mathsf{tid}' = \mathsf{tid}$ and for each $\mathsf{tid}' \in T$: let $\mathsf{h} \in \mathsf{HEAD}$, $\mathsf{h} \geq \mathsf{h}_3$ ($\mathsf{h} := \mathsf{h}_3$ for $\mathsf{tid}' = \mathsf{tid}$), $\mathsf{last\text{-}key}(\mathsf{tid}', \mathsf{a}, \mathsf{h}) < \mathsf{k} \leq \mathsf{last\text{-}key}_g(\mathsf{tid}', \mathsf{a}, \mathsf{h}{+}1)$, then $\mathsf{mem\text{-}value}' := \mathsf{mem\text{-}value}'[(\mathsf{tid}', \mathsf{a}, \mathsf{h}) \hookleftarrow \mathsf{v}]$, $\mathsf{last\text{-}key}' := \mathsf{last\text{-}key}'[(\mathsf{tid}', \mathsf{a}, \mathsf{h}) \hookleftarrow \mathsf{k}]$, $\lambda := \lambda \cdot (\mathsf{h}, \mathsf{prop}, \mathsf{tid}', \mathsf{tid}, i, \mathsf{a})$.

For brevity we allowed a single transition to be labeled by several events. An automaton with such transitions can be trivially translated to the canonical form by breaking one such transition into several consecutive ones.

Final states The set of final states F_M is a subset of $S_M \setminus \{s_{0M}\}$ consisting of all states with $\mathsf{mem\text{-}value}(\mathsf{tid}, \mathsf{a}, \mathsf{h}) = \mathsf{mem\text{-}value}_g(\mathsf{tid}, \mathsf{a}, \mathsf{h}{+}1)$, $\mathsf{last\text{-}key}(\mathsf{tid}, \mathsf{a}, \mathsf{h}) = \mathsf{last\text{-}key}_g(\mathsf{tid}, \mathsf{a}, \mathsf{h}+1)$ for all $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{a} \in \mathsf{ADDR}$, $\mathsf{h} \in [1..n-1]$.

Soundness and completeness

Lemma 16. $\mathcal{L}(M) \subseteq \mathcal{C}_{\mathit{power}}(\mathcal{P})$.

Proof. Consider $\sigma = \lambda_1 \cdots \lambda_m$, such that $s_{0M} \xrightarrow{\lambda_1} s_{M1} \xrightarrow{\lambda_2} \cdots \xrightarrow{\lambda_m} s_{Mm} \in F_M$. For $\mathsf{h} \in \mathsf{HEAD}$, let $\tau^s_\mathsf{h} := \mathsf{second}((\lambda_1 \cdots \lambda_s)\downarrow(\{\mathsf{h}\} \times \mathsf{E}))$, $s \in [0..m]$.

Let $(s_{Z}{}^0_1 \ldots s_{Z}{}^0_n) \in (S_Z)^n$ be the states of Z defined so that SND-B holds for $s = 0$ (see below). By induction on $s \in [1..m]$ we show:

SND-A $s_{Z}{}^0_\mathsf{h} \xrightarrow{\tau^s_\mathsf{h}} s_{Z}{}^s_\mathsf{h}$.

SND-B For all $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{h} \in \mathsf{HEAD}$, $s_{Z}{}^s_\mathsf{h} = (\mathsf{ts}, (\mathsf{co}, \mathsf{prop}))$, $\mathsf{ts}(\mathsf{tid}) = (\mathsf{fetched}, \mathsf{committed}, \mathsf{loaded})$ holds:

SND-B1 $\mathsf{fetched}$ is the list of instructions fetched by $(\tau_1^m \cdots \tau_{\mathsf{h}-1}^m \cdot \tau_{\mathsf{h}}^s)\downarrow \mathsf{fetch}\downarrow\mathsf{tid}$.

SND-B2 $\mathsf{committed}$ consists of the indices of instructions committed by $(\tau_1^m \cdots \tau_{\mathsf{h}-1}^m \cdot \tau_{\mathsf{h}}^s)\downarrow\mathsf{commit}\downarrow\mathsf{tid}$.

SND-B3 loaded contains the information about the stores being read by loads in $(\tau_1^m \cdots \tau_{\mathsf{h}-1}^m \cdot \tau_{\mathsf{h}}^s)$ determined according to Lemmas 4 and 5.

SND-B4 $\mathsf{co}(\mathsf{tid}, i) = \mathsf{k}$ if $(\mathsf{commit}, \mathsf{tid}, i, \mathsf{k}, \mathsf{a}) \in \tau_1^m \cdots \tau_{\mathsf{h}-1}^m \cdot \tau_{\mathsf{h}}^s$ for some $\mathsf{a} \in \mathsf{ADDR}$, otherwise, $\mathsf{co}(\mathsf{tid}, i) = \bot$.

SND-B5 $\mathsf{prop}(\mathsf{tid}, \mathsf{a}) = (\mathsf{tid}', i')$ if $(\mathsf{prop}, \mathsf{tid}, \mathsf{tid}', i', \mathsf{a}) = \mathsf{last}((\tau_1^m \cdots \tau_{\mathsf{h}-1}^m \cdot \tau_{\mathsf{h}}^s)\downarrow(\mathsf{prop}, \mathsf{tid}, *, *, \mathsf{a}))$, otherwise, $\mathsf{prop}(\mathsf{tid}, \mathsf{a}) = \mathsf{init}_{\mathsf{a}}$.

SND-C For each $\mathsf{tid} \in \mathsf{TID}$: $\mathsf{ctrl\text{-}state}(\mathsf{tid}) = \mathsf{dst}(\mathsf{last}(s_{Z_1^s}.\mathsf{ts}(\mathsf{tid}).\mathsf{fetched}))$ (or $q_{0_{\mathsf{tid}}}$ if no instructions were fetched).

SND-D For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{r} \in \mathsf{REG}$, for each $\mathsf{h} \in [\mathsf{reg\text{-}comp\text{-}head}(\mathsf{tid}, \mathsf{r})..n]$: $\mathsf{reg\text{-}value}(\mathsf{tid}, \mathsf{r}) = \mathsf{eval}(\mathsf{tid}, \mathsf{instr\text{-}count}(\mathsf{tid}) + 1, \mathsf{r})$ computed for the state $s_{Z_{\mathsf{h}}^s}$.

SND-E For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{r} \in \mathsf{REG}$, $\mathsf{h} \in [\mathsf{reg\text{-}comm\text{-}head}(\mathsf{tid}, \mathsf{r})..n]$: let i be the index of the latest instruction in $s_{Z_{\mathsf{h}}^s}.\mathsf{ts}(\mathsf{tid}).\mathsf{fetched}$ writing to r, then $i \in s_{Z_{\mathsf{h}}^s}.\mathsf{ts}(\mathsf{tid}).\mathsf{committed}$.

SND-F For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{h} \in [\mathsf{assume\text{-}comm\text{-}head}(\mathsf{tid})..n]$: $s_{Z_{\mathsf{h}}^s}$ does not contain uncommitted conditional instructions in thread tid having indices $\leq \mathsf{instr\text{-}count}(\mathsf{tid})$.

SND-G For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{a} \in \mathsf{ADDR}$, $\mathsf{h} \in \mathsf{HEAD}$: let $w := s_{Z_{\mathsf{h}}^s}.\mathsf{prop}(\mathsf{tid}, \mathsf{a})$. If $w = \mathsf{init}_{\mathsf{a}}$, $\mathsf{mem\text{-}value}(\mathsf{tid}, \mathsf{a}, \mathsf{h}) = 0$. If $w = (\mathsf{tid}', i')$, $\mathsf{mem\text{-}value}(\mathsf{tid}, \mathsf{a}, \mathsf{h}) = \mathsf{val}(\mathsf{tid}', i')$ computed in $s_{Z_{\mathsf{h}}^s}$.

SND-H For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{a} \in \mathsf{ADDR}$, $\mathsf{h} \in \mathsf{HEAD}$: $\mathsf{last\text{-}key}_g(\mathsf{tid}, \mathsf{a}, \mathsf{h}) \leq s_{Z_{\mathsf{h}}^s}.\mathsf{co}(s_{Z_{\mathsf{h}}^s}.\mathsf{prop}(\mathsf{tid}, \mathsf{a})) = \mathsf{last\text{-}key}(\mathsf{tid}, \mathsf{a}, \mathsf{h}) \leq \mathsf{last\text{-}key}_g(\mathsf{tid}, \mathsf{a}, \mathsf{h} + 1)$.

SND-K For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{a} \in \mathsf{ADDR}$, $\mathsf{h} \in \mathsf{HEAD}$: let $i \in \mathbb{N}$ be the maximal index, such that $s_{Z_{\mathsf{h}}^s}.\mathsf{ts}(\mathsf{tid}).\mathsf{fetched}[i]$ is a store, $\mathsf{addr}(\mathsf{tid}, i) = \mathsf{a}$ in $s_{Z_n^s}$. Let i' be the maximal index, such that $s_{Z_{\mathsf{h}}^s}.\mathsf{ts}(\mathsf{tid}).\mathsf{fetched}[i']$ is a store, $\mathsf{addr}(\mathsf{tid}, i') \in \{\bot, \mathsf{a}\}$ in $s_{Z_{\mathsf{h}}^s}$. Then $\mathsf{early\text{-}mem\text{-}value}(\mathsf{tid}, i, \mathsf{h}) = \bot$ if such i does not exist or $i \in s_{Z_{\mathsf{h}}^s}.\mathsf{ts}(\mathsf{tid}).\mathsf{committed}$. Otherwise, $\mathsf{early\text{-}mem\text{-}value}(\mathsf{tid}, i, \mathsf{h}) = \top$ if $\mathsf{addr}(\mathsf{tid}, i') = \bot$ or $\mathsf{val}(\mathsf{tid}, i) = \bot$ in $s_{Z_{\mathsf{h}}^s}$. Otherwise, $\mathsf{early\text{-}mem\text{-}value}(\mathsf{tid}, i, \mathsf{h}) = \mathsf{val}(\mathsf{tid}, i)$ computed in $s_{Z_{\mathsf{h}}^s}$.

SND-L For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{h} \in [\mathsf{addr\text{-}comp\text{-}head}(\mathsf{tid})..n]$, $i \in [1..|s_{Z_{\mathsf{h}}^s}.\mathsf{ts}(\mathsf{tid}).\mathsf{fetched}|]$: $\mathsf{addr}(\mathsf{tid}, i) \neq \bot$ in $s_{Z_{\mathsf{h}}^s}$.

SND-M For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{a} \in \mathsf{ADDR}$, $\mathsf{h} \in [\mathsf{addr\text{-}comm\text{-}head}(\mathsf{tid}, \mathsf{a})..n]$: if $\mathsf{addr}(\mathsf{tid}, i) = \mathsf{a}$ in $s_{Z_n^s}$ for some i, then $i \in s_{Z_{\mathsf{h}}^s}.\mathsf{ts}(\mathsf{tid}).\mathsf{committed}$.

Finally we will show that $s_{Z_{\mathsf{h}}^m} = s_{Z_{\mathsf{h}+1}^0}$ for all $\mathsf{h} \in [1..n-1]$ and $s_{Z_n^m} \in F_Z$, thus proving the claim of the lemma.

Base case: $s = 1$, we must show that there s_{M_1} satisfies the inductive statement. This is easy to check by definition of the destination state of MH-GUESS transition.

Step case: assume the inductive statement holds for some $s \in [0..m-1]$. Consider λ_s (for notational convenience and without loss of generality we assume below that $\mathsf{h}_j \neq \mathsf{h}_{j'}$ for $j \neq j'$):

Assignment $\lambda_s = (\mathsf{h}_1, \mathsf{fetch}, \mathsf{tid}, \mathsf{instr}) \cdot (\mathsf{h}_3, \mathsf{commit}, \mathsf{tid}, i)$, $\mathsf{instr} = q_1 \xrightarrow{\mathsf{r} \leftarrow e_\mathsf{v}} q_2$. Let $\mathsf{e}_1 := (\mathsf{fetch}, \mathsf{tid}, \mathsf{instr})$, $\mathsf{e}_3 := (\mathsf{commit}, \mathsf{tid}, i)$.

We need to show that $s_{Z^{s-1}_{\mathsf{h}_1}} \xrightarrow{\mathsf{e}_1} s_{Z^{s}_{\mathsf{h}_1}}$, i.e. that the assignment instruction can be fetched. This follows from the choice of $\mathsf{h}_1 := 1$ in MH-ASSIGN and SND-B1, SND-C.

We also need to show that $s_{Z^{s-1}_{\mathsf{h}_3}} \xrightarrow{\mathsf{e}_3} s_{Z^{s}_{\mathsf{h}_3}}$, i.e. that the assignment instruction can be committed. First, the e_3 transition requires the instruction being committed to be fetched, which holds due to SND-B1 and $\mathsf{h}_3 \geq \mathsf{h}_1$. Second, this instruction must be not committed yet, which holds by SND-B2 and the fact that M commits each instruction once and only once. Third, all control dependencies must be committed. This is by the choice of h_3 in MH-ASSIGN and SND-F. Fourth, all the preceding data dependencies must be committed. This is by the choice of h_3 in MH-ASSIGN and SND-E. Finally, the argument of the function must be computed. This is by choice of $\mathsf{h}_3 \geq \mathsf{h}_2$ in MH-ASSIGN, Lemma 3, and SND-D.

In the end, we must show that the invariants hold in the new state. The only non-trivial thing is SND-D, which holds due to SND-D, definition of v in MH-ASSIGN, definitions of eval, and the fact that functions in FUN are deterministic.

Assume $\lambda_s = (\mathsf{h}_1, \mathsf{fetch}, \mathsf{tid}, q_1 \xrightarrow{\mathsf{instr}} q_2) \cdot (\mathsf{h}_3, \mathsf{commit}, \mathsf{tid}, i)$, $\mathsf{instr} = \mathtt{assume}(e_\mathsf{v})$. The proof is similar to the previous case. The commit transition additionally requires $\mathsf{eval}(\mathsf{tid}, i, e_\mathsf{v}) \neq 0$, which holds due to the fact that a similar check in MH-ASSUME holds, SND-D, definitions of eval, the fact that functions in FUN are deterministic.

Load $\lambda_s = (\mathsf{h}_1, \mathsf{fetch}, \mathsf{tid}, \mathsf{instr}) \cdot (\mathsf{h}_2, \mathsf{load}, \mathsf{tid}, i, \mathsf{a}) \cdot (\mathsf{h}_3, \mathsf{commit}, \mathsf{tid}, i)$, $\mathsf{instr} = \mathsf{r} \leftarrow \mathtt{mem}[e_\mathsf{a}]$. Let $\mathsf{e}_1 := (\mathsf{fetch}, \mathsf{tid}, \mathsf{instr})$, $\mathsf{e}_2 := (\mathsf{load}, \mathsf{tid}, i, \mathsf{a})$, $\mathsf{e}_3 := (\mathsf{commit}, \mathsf{tid}, i)$.

$s_{Z^{s-1}_{\mathsf{h}_1}} \xrightarrow{\mathsf{e}_1} s_{Z^{s}_{\mathsf{h}_1}}$ holds for the same reasons as before.

Next, we show that $s_{Z^{s-1}_{\mathsf{h}_2}} \xrightarrow{\mathsf{e}_2} s_{Z^{s}_{\mathsf{h}_2}}$, where this transition is a POW-EARLY transition in the early read case of MH-LOAD and a POW-LOAD transition in the load from memory case. First, we must show that $s_{Z^{s}_{\mathsf{h}}}.\mathsf{ts}(\mathsf{tid}).\mathsf{loaded}[i] = \bot$. This holds by SND-B3 and the fact that M generates a load event once and only once for a single fetched load instruction.

Assume the early read case. This means, $\mathsf{early\text{-}mem\text{-}value}(\mathsf{tid}, \mathsf{a}, \mathsf{h}_2) \in \mathsf{DOM}$. By SND-K, this means, the last fetched store with an unknown address or address of the load is not yet committed, has the address of the load and has the value known. By POW-EARLY, the load can take the value from this store, and SND-B3 holds in the new state.

Consider the load from memory case. This means, $\mathsf{early\text{-}mem\text{-}value}(\mathsf{tid}, \mathsf{a}, \mathsf{h}_2) = \bot$. By SND-K, this means, there is no earlier fetched store with the same address which is not yet committed. By POW-LOAD, the load can take the value from the last propagated store, and SND-B3 holds in the new state.

Argumentation for $s_{Z^{s-1}_{\mathsf{h}_3}} \xrightarrow{\mathsf{e}_3} s_{Z^{s}_{\mathsf{h}_3}}$ is similar to the previous cases. Additionally, first we must show that $s_{Z^{s}_{\mathsf{h}}}.\mathsf{ts}(\mathsf{tid}).\mathsf{loaded}[i] \neq \bot$. This is by $\mathsf{h}_3 \geq \mathsf{h}_2$

(MH-LOAD), SND-B3. Second, we must ensure that all preceding instructions accessing the same address a are committed, and there are no previously fetched instructions with unknown address. This holds by choice of h_3 in MH-LOAD, SND-L, and SND-M.

In the new state, SND-D holds by definition of v in POW-LOAD, definitions of eval, SND-G, and SND-K. Proofs for the other conditions are simpler.

Store $\lambda_s = (\mathsf{h}_1, \mathsf{fetch}, \mathsf{tid}, \mathsf{instr}) \cdot (\mathsf{h}_3, \mathsf{commit}, \mathsf{tid}, i, \mathsf{k}, \mathsf{a}) \cdot (\mathsf{h}_3, \mathsf{prop}, \mathsf{tid}, \mathsf{tid}, i, \mathsf{a}) \cdot (\mathsf{h}_4, \mathsf{prop}, \mathsf{tid}_1, \mathsf{tid}, i, \mathsf{a}) \cdots (\mathsf{h}_{u+3}, \mathsf{prop}, \mathsf{tid}_u, \mathsf{tid}, i, \mathsf{a})$. Let $\mathsf{e}_1 := (\mathsf{fetch}, \mathsf{tid}, \mathsf{instr})$, $\mathsf{e}_3 := (\mathsf{commit}, \mathsf{tid}, i, \mathsf{k}, \mathsf{a})$, $\mathsf{e}_4 := (\mathsf{prop}, \mathsf{tid}, \mathsf{tid}, i, \mathsf{a})$, $\mathsf{e}_{j+3} := (\mathsf{prop}, \mathsf{tid}_j, \mathsf{tid}, i, \mathsf{a})$ for $j \in [1..u]$.

${s_Z}_{\mathsf{h}_1}^{s-1} \xrightarrow{\mathsf{e}_1} {s_Z}_{\mathsf{h}_1}^{s}$ holds for the same reasons as before.

${s_Z}_{\mathsf{h}_3}^{s-1} \xrightarrow{\mathsf{e}_3} {s_Z}_{\mathsf{h}_3}^{s}$ holds for the same reasons as in the case of a load. The requirement that the coherence key is unique in POW-STORE follows from a similar requirement in MH-STORE and SND-H. By POW-STORE, the only available transition from ${s_Z}_{\mathsf{h}_2}^{s}$ is a propagation of the write to its thread, i.e. e_4, which indeed follows e_3 in τ. Next, we show that e_4 and further propagate transitions are feasible.

First, POW-PROP rule requires the write being propagated to have a coherence key (i.e. to be committed), which holds by choice of h_j, $j \in [3..u+3]$ in MH-STORE and SND-B2. Second, it requires the coherence key of the latest propagated store to be less than the key of the store being propagated. This is adhered due to the check $\mathsf{last\text{-}key}(\mathsf{tid}', \mathsf{a}, \mathsf{h}) < \mathsf{k}$ and SND-H.

It is easy to see that the inductive statements hold in the new state as well.

Now we prove ${s_Z}_{\mathsf{h}}^{m} = {s_Z}_{\mathsf{h}+1}^{0}$ for all $\mathsf{h} \in [1..n-1]$. The equality of ts components immediately follows from SND-B inductive statement.

Now we prove ${s_Z}_{n}^{m} \in F_Z$. FIN-COMM holds, because Z always emits a commit event for each fetched instruction.

Let us turn to FIN-LD property. First, one should note that M generates prop events for stores to the same address in each part τ_j in the ascending order by k. This is by MH-STORE. Together with SND-H, this means that these events are sorted in τ in the ascending order by k. The rest of the proof of FIN-LD is a simple case consideration: whether the loads i, i' were done from memory or from a local store early.

FIN-LD-ST is proven by a similar case consideration.

□

We call α a *prefix* of σ and write $\alpha \sqsubseteq \sigma$ if $\sigma = \alpha \cdot \beta$ for some β.

Lemma 17. *$\{\tau \in \mathsf{C}_{\mathit{power}}(\mathcal{P}) \mid \tau$ is in normal form of degree $n\} \subseteq \mathcal{L}(M)$.*

Proof. Let $\tau = \tau_1 \cdots \tau_n \in \mathsf{C}_{\mathsf{mm}}(\mathcal{P})$ be a normal-form computation, i.e. $s_{0Z} \xrightarrow{\tau} s_Z \in F_Z$. We show that there is a sequence of transitions $s_{0M} \xrightarrow{\lambda_1} s_{M1} \xrightarrow{\lambda_2} \ldots \xrightarrow{\lambda_m} s_{Mm} \in F_M$, such that $\tau_{\mathsf{h}} = \mathsf{second}((\lambda_1 \cdots \lambda_n)\downarrow(\{\mathsf{h}\} \times \mathsf{E}))$.

Let $\tau_{\mathsf{h}}^{s} := \mathsf{second}((\lambda_1 \cdots \lambda_s) \downarrow (\{\mathsf{h}\} \times \mathsf{E}))$, $s_{0Z} \xrightarrow{\tau_1 \cdots \tau_{\mathsf{h}-1} \cdot \tau_{\mathsf{h}}^{s}} {s_Z}_{\mathsf{h}}^{s} \to^* s_Z$. By induction on $s \in [1, \infty)$ we show the following inductive statements:

CMPL-A There is a sequence of s transitions: $s_{0M} \xrightarrow{\lambda_1} s_{M1} \xrightarrow{\lambda_2} \ldots \xrightarrow{\lambda_s} s_{Ms}$.

CMPL-B For all $\mathsf{h} \in \mathsf{HEAD}$: $\tau_{\mathsf{h}} = \tau^s_{\mathsf{h}}.\overline{\tau^s_{\mathsf{h}}}$ for some $\overline{\tau^s_{\mathsf{h}}}$.

CMPL-C If $\mathsf{e}_1, \mathsf{e}_2 \in \tau$ are two events related to instruction (tid, i), then $\mathsf{e}_1 \in \tau^s_{\mathsf{h}}$ for some h iff $\mathsf{e}_2 \in \tau^s_{\mathsf{h}'}$ for some h'.

CMPL-D For each $\mathsf{tid} \in \mathsf{TID}$: $\mathsf{ctrl\text{-}state}(\mathsf{tid}) = \mathsf{dst}(\mathsf{last}(s_{Z1}.\mathsf{ts}(\mathsf{tid}).\mathsf{fetched}))$ (or $\mathsf{ctrl\text{-}state}(\mathsf{tid}) = q_{0\mathsf{tid}}$ if no instructions were fetched).

CMPL-F For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{r} \in \mathsf{REG}$, $\mathsf{h} \in [\mathsf{reg\text{-}comp\text{-}head}(\mathsf{tid}, \mathsf{r})..n]$: $\mathsf{reg\text{-}value}(\mathsf{tid}, \mathsf{r}) = \mathsf{eval}(\mathsf{tid}, \mathsf{instr\text{-}count}(\mathsf{tid}) + 1, \mathsf{r})$ computed for the state $s_{Z\mathsf{h}}$.

CMPL-F' For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{r} \in \mathsf{REG}$, $\mathsf{h} \in [1..\mathsf{reg\text{-}comp\text{-}head}(\mathsf{tid}, \mathsf{r}) - 1]$: $\mathsf{eval}(\mathsf{tid}, \mathsf{instr\text{-}count}(\mathsf{tid}) + 1, \mathsf{r}) = \bot$.

CMPL-G For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{r} \in \mathsf{REG}$, $\mathsf{h} \in [\mathsf{reg\text{-}comm\text{-}head}(\mathsf{tid})..n]$: let i be the index of the last instruction in $s_{Z\mathsf{h}}.\mathsf{ts}(\mathsf{tid}).\mathsf{fetched}$ writing to r, then $i \in s_{Z\mathsf{h}}.\mathsf{ts}(\mathsf{tid}).\mathsf{committed}$.

CMPL-G' For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{r} \in \mathsf{REG}$, $\mathsf{h} \in [1..\mathsf{reg\text{-}comm\text{-}head}(\mathsf{tid}, \mathsf{r}) - 1]$: let i be the index of the last instruction in $s_{Z\mathsf{h}}.\mathsf{ts}(\mathsf{tid}).\mathsf{fetched}$, then $i \notin s_{Z\mathsf{h}}.\mathsf{ts}(\mathsf{tid}).\mathsf{committed}$.

CMPL-K For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{h} \in [\mathsf{assume\text{-}comm\text{-}head}(\mathsf{tid})..n]$: let i be an index of an `assume()` instruction in $s_{Z\mathsf{h}}.\mathsf{ts}(\mathsf{tid}).\mathsf{fetched}$, then $i \in s_{Z\mathsf{h}}.\mathsf{ts}(\mathsf{tid}).\mathsf{committed}$.

CMPL-K' For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{h} \in [1..\mathsf{assume\text{-}comm\text{-}head}(\mathsf{tid}) - 1]$: let i be an index of the last `assume()` instruction in $s_{Z\mathsf{h}}.\mathsf{ts}(\mathsf{tid}).\mathsf{fetched}$, then $i \notin s_{Z\mathsf{h}}.\mathsf{ts}(\mathsf{tid}).\mathsf{committed}$.

CMPL-L For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{a} \in \mathsf{ADDR}$, $\mathsf{h} \in \mathsf{HEAD}$: let $w := s_{Z\mathsf{h}}.\mathsf{prop}(\mathsf{tid}, \mathsf{a})$. If $w = \mathsf{init}_{\mathsf{a}}$, $\mathsf{mem\text{-}value}(\mathsf{tid}, \mathsf{a}, \mathsf{h}) = 0$. If $w = (\mathsf{tid}', i')$, $\mathsf{mem\text{-}value}(\mathsf{tid}, \mathsf{a}, \mathsf{h}) = \mathsf{val}(\mathsf{tid}', i')$ computed in $s_{Z\mathsf{h}}$.

CMPL-M For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{a} \in \mathsf{ADDR}$, $\mathsf{h} \in \mathsf{HEAD}$: $\mathsf{last\text{-}key}_g(\mathsf{tid}, \mathsf{a}, \mathsf{h}) < s_{Z\mathsf{h}}.\mathsf{co}(s_{Z\mathsf{h}}.\mathsf{prop}(\mathsf{tid}, \mathsf{a})) = \mathsf{last\text{-}key}(\mathsf{tid}, \mathsf{a}, \mathsf{h}) \leq \mathsf{last\text{-}key}_g(\mathsf{tid}, \mathsf{a}, \mathsf{h} + 1)$.

CMPL-N For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{a} \in \mathsf{ADDR}$, $\mathsf{h} \in \mathsf{HEAD}$: let $i \in \mathbb{N}$ be the maximal index, such that $s_{Z\mathsf{h}}.\mathsf{ts}(\mathsf{tid}).\mathsf{fetched}[i]$ is a store, $\mathsf{addr}(\mathsf{tid}, i) = \mathsf{a}$ in s_{Zn}. Let i' be the maximal index, such that $s_{Z\mathsf{h}}.\mathsf{ts}(\mathsf{tid}).\mathsf{fetched}[i']$ is a store, $\mathsf{addr}(\mathsf{tid}, i') \in \{\bot, \mathsf{a}\}$ in $s_{Z\mathsf{h}}$. Then $\mathsf{early\text{-}mem\text{-}value}(\mathsf{tid}, i, \mathsf{h}) = \bot$ if such i does not exist or $i \in s_{Z\mathsf{h}}.\mathsf{ts}(\mathsf{tid}).\mathsf{committed}$. Otherwise, $\mathsf{early\text{-}mem\text{-}value}(\mathsf{tid}, i, \mathsf{h}) = \top$ if $\mathsf{addr}(\mathsf{tid}, i') = \bot$ or $\mathsf{val}(\mathsf{tid}, i) = \bot$ in $s_{Z\mathsf{h}}$. Otherwise, $\mathsf{early\text{-}mem\text{-}value}(\mathsf{tid}, i, \mathsf{h}) = \mathsf{val}(\mathsf{tid}, i)$ computed in $s_{Z\mathsf{h}}$.

CMPL-P For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{a} \in \mathsf{ADDR}$, $\mathsf{h} \in [\mathsf{addr\text{-}comm\text{-}head}(\mathsf{tid}, \mathsf{a})..n]$: if $\mathsf{addr}(\mathsf{tid}, i) = \mathsf{a}$ in s_{Zn} for some i, then $i \in s_{Z\mathsf{h}}.\mathsf{ts}(\mathsf{tid}).\mathsf{committed}$.

CMPL-P' For each $\mathsf{tid} \in \mathsf{TID}$, $\mathsf{a} \in \mathsf{ADDR}$, $\mathsf{h} \in [1..\mathsf{addr\text{-}comm\text{-}head}(\mathsf{tid}, \mathsf{a}) - 1]$: there is i with $\mathsf{addr}(\mathsf{tid}, i) = \mathsf{a}$ in s_{Zn}, such that $i \in s_{Z\mathsf{h}}.\mathsf{ts}(\mathsf{tid}).\mathsf{committed}$.

CMPL-R For each $\mathsf{tid} \in \mathsf{TID}$: $\mathsf{instr\text{-}count}(\mathsf{tid}) = |s_{Z1}.\mathsf{ts}(\mathsf{tid}).\mathsf{fetched}|$.

Base case: $s = 1$. We choose the first (MH-GUESS) transition $s_{0M} \xrightarrow{\lambda_1} s_{M1}$, so that the inductive statements hold:

Guess We define $\mathsf{mem\text{-}value}$ and $\mathsf{last\text{-}key}$ components of s_{M1} according to CMPL-L and CMPL-M requirements. The other inductive statements trivially hold.

Assume the inductive statements hold for s and $\overline{\tau^s_{\mathsf{h}}} \neq \varepsilon$ for some $\mathsf{h} \in \mathsf{HEAD}$. We show they hold for $s' := s + 1$. The proof is done by pointing out an appropriate transition $s_{Ms} \xrightarrow{\lambda_{s+1}} s_{Ms+1}$. We choose the first possible option out of the following:

Assignment Assume $\mathsf{e}_1 \sqsubseteq \overline{\tau^s_{\mathsf{h}_1}}$, $\mathsf{e}_3 \sqsubseteq \overline{\tau^s_{\mathsf{h}_3}}$, where $\mathsf{h}_1 < \mathsf{h}_3$ ($\mathsf{h}_1 = \mathsf{h}_3$ is possible, but here and further we write strict inequalities for notational convenience), $\mathsf{e}_1 := (\mathsf{fetch}, \mathsf{tid}, q_1 \xrightarrow{\mathsf{cmd}} q_2)$, $\mathsf{e}_3 := (\mathsf{commit}, \mathsf{tid}, i)$, $\mathsf{h}_1 = 1$, $i = \mathsf{instr\text{-}count}(\mathsf{tid})$, $\mathsf{cmd} = \mathsf{r} \leftarrow e_{\mathsf{v}}$. Then, as we show next, a MH-ASSIGN transition is feasible.
First, $s_{Z^s_{\mathsf{h}_1}} \xrightarrow{\mathsf{e}_1}$, therefore, the state of the last fetched instruction in thread tid in $s_{Z^s_{\mathsf{h}_1}}$ is q_1. By CMPL-D, $\mathsf{ctrl\text{-}state}(\mathsf{tid}) = q_1$ too.
Second, we choose $\mathsf{h}_2 := \max\{\mathsf{reg\text{-}comm\text{-}head}(\mathsf{tid}, \mathsf{r}) \mid \mathsf{r} \text{ is read in } \mathsf{cmd}\}$. It satisfies the requirements from MH-ASSIGN. Note that $\mathsf{h}_2 \leq \mathsf{h}_3$ by CMPL-F' and POW-COMMIT: an instruction cannot be committed, until its arguments are computed.
Third, we must show that for each register r read by the instruction holds $\mathsf{h}_3 \geq \mathsf{reg\text{-}comm\text{-}head}(\mathsf{tid}, \mathsf{r})$ and h_3. This holds by CMPL-G', CMPL-K', and POW-COMMIT: an instruction cannot be committed until its data and control dependencies are committed.
In the destination state, CMPL-F holds by CMPL-F in the source state, definition of $\mathsf{reg\text{-}value}'$ in MH-ASSIGN and definitions of eval. The other inductive statements trivially hold.

Assume Assume $\mathsf{e}_1 \sqsubseteq \overline{\tau^s_{\mathsf{h}_1}}$, $\mathsf{e}_3 \sqsubseteq \overline{\tau^s_{\mathsf{h}_3}}$, where $\mathsf{h}_1 < \mathsf{h}_3$, $\mathsf{e}_1 = (\mathsf{fetch}, \mathsf{tid}, q_1 \xrightarrow{\mathsf{cmd}} q_2)$, $\mathsf{e}_3 = (\mathsf{commit}, \mathsf{tid}, i)$, where $i = \mathsf{instr\text{-}count}(\mathsf{tid})$, $\mathsf{h}_1 = 1$, $i = \mathsf{instr\text{-}count}(\mathsf{tid})$, $\mathsf{cmd} = \mathtt{assume}(e_{\mathsf{v}})$. Then, a MH-ASSUME transition is feasible.
The proof is similar to the proof for the case of assignment. The MH-ASSUME transition additionally requires $\mathsf{eval}(\mathsf{tid}, e_{\mathsf{v}}) \neq 0$. This holds by CMPL-F, definition of $\mathsf{reg\text{-}value}'$ in MH-ASSIGN and definitions of eval.
The inductive statements trivially hold in the destination state.

Load Assume $\mathsf{e}_1 \sqsubseteq \overline{\tau^s_{\mathsf{h}_1}}$, $\mathsf{e}_2 \sqsubseteq \overline{\tau^s_{\mathsf{h}_2}}$, $\mathsf{e}_3 \sqsubseteq \overline{\tau^s_{\mathsf{h}_3}}$, where $\mathsf{h}_1 < \mathsf{h}_2 < \mathsf{h}_3$, $\mathsf{e}_1 = (\mathsf{fetch}, \mathsf{tid}, q_1 \xrightarrow{\mathsf{cmd}} q_2)$, $\mathsf{e}_2 = (\mathsf{load}, \mathsf{tid}, i, \mathsf{a})$, $\mathsf{e}_3 = (\mathsf{commit}, \mathsf{tid}, i)$, $i = \mathsf{instr\text{-}count}(\mathsf{tid})$, $\mathsf{cmd} = \mathsf{r} \leftarrow \mathtt{mem}[e_{\mathsf{v}}]$. We show that a MH-LOAD transition is feasible. We point out only differences with respect to the proof for the assignment case.
Assume e_2 was produced by a POW-EARLY transition. This means, the last store writing to a has its address known and is not committed yet in $s_{Z^s_{\mathsf{h}_2}}$. Then, by CMPL-N, $\mathsf{early\text{-}mem\text{-}value}(\mathsf{tid}, \mathsf{a}, \mathsf{h}_2) \in \mathsf{DOM}$, and we have $\mathsf{v} := \mathsf{early\text{-}mem\text{-}value}(\mathsf{tid}, \mathsf{a}, \mathsf{h}_2)$. Assume e_2 was produced by a POW-LOAD transition. Then, POW-EARLY transition was not possible (Lemma 4, Lemma 5). This means, there was no in-flight stores to a in $s_{Z^s_{\mathsf{h}_2}}$. Then, by CMPL-N, $\mathsf{early\text{-}mem\text{-}value}(\mathsf{tid}, \mathsf{a}, \mathsf{h}_2) = \bot$, and we have $\mathsf{v} := \mathsf{mem\text{-}value}(\mathsf{tid}, \mathsf{a}, \mathsf{h}_2)$. In both cases, by CMPL-N, CMPL-L we have $\mathsf{reg\text{-}value}'$ and $\mathsf{reg\text{-}comp\text{-}head}'$ satisfying CMPL-F and CMPL-F'.
Additionally, we must show that $\mathsf{h}_3 \geq \mathsf{addr\text{-}comm\text{-}head}(\mathsf{tid}, \mathsf{a})$. This holds by CMPL-P' and CMPL-N.

Store Assume $u \in \mathbb{N}$, $\mathsf{e}_j \sqsubseteq \overline{\tau^s_{\mathsf{h}_j}}$ for $j \in [1..u+3]$, where $\mathsf{h}_2 = \mathsf{h}_3$, $\mathsf{e}_1 = (\mathsf{fetch}, \mathsf{tid}, q_1 \xrightarrow{\mathsf{cmd}} q_2)$, $\mathsf{e}_2 = (\mathsf{commit}, \mathsf{tid}, i, \mathsf{k}, \mathsf{a})$, $\mathsf{e}_3 = (\mathsf{prop}, \mathsf{tid}, \mathsf{tid}, i, \mathsf{a})$, $\mathsf{e}_j = (\mathsf{prop}, \mathsf{tid}_j, \mathsf{tid}, i, \mathsf{a})$ for $j \in [4..u+3]$, $i = \mathsf{instr\text{-}count}(\mathsf{tid})$, $\mathsf{cmd} = \mathtt{mem}[e_\mathsf{a}] \leftarrow e_\mathsf{v}$. Assume that there are no other prop events for (tid, i) in τ, except for $\mathsf{e}_3 \dots \mathsf{e}_{u+3}$. We show that a MH-STORE transition is feasible. The requirements to be checked are similar to those in the load case. The requirement that k is not already used holds by CMPL-M and the fact that the same requirement in POW-STORE is met.
Consider the requirements in MH-STORE for generating prop events. The requirement that propagation event to thread tid is generated in the same part as commit is met by assumption $\mathsf{h}_3 = \mathsf{h}_2$. The requirement $\mathsf{last\text{-}key}(\mathsf{tid}', \mathsf{a}, \mathsf{h}) < \mathsf{k} \leq \mathsf{last\text{-}key}_g(\mathsf{tid}', \mathsf{a}, \mathsf{h}+1)$ is met by CMPL-L, choice of $\mathsf{last\text{-}key}_g$ in the initial transition, and POW-PROP.
This means, inductive invariant CMPL-A holds for $s+1$. Also, CMPL-B holds by choice of $\mathsf{e}_1 \dots \mathsf{e}_{u+3}$, CMPL-D holds trivially. CMPL-C holds by assumption that there are no other prop events in τ, except for $\mathsf{e}_3 \dots \mathsf{e}_{u+3}$. CMPL-F, CMPL-F', CMPL-G, CMPL-G' hold as store instruction does not affect register values. CMPL-K, CMPL-K' hold as a store instruction is not `assume()`. CMPL-L holds by definition of $\mathsf{mem\text{-}value}'$ in MH-STORE. CMPL-M holds by definition of $\mathsf{last\text{-}key}'$ in MH-STORE. CMPL-N holds by definition of $\mathsf{early\text{-}mem\text{-}value}'$ in MH-STORE. CMPL-P, CMPL-P' hold by definition of $\mathsf{addr\text{-}comm\text{-}head}'$ in MH-STORE. CMPL-R hold by definition of $\mathsf{instr\text{-}count}'$ in MH-STORE.

Now we must show that one of the cases above always takes place. Consider the event $\mathsf{e} = \mathsf{first}(\overline{\tau^s_1})$. By CMPL-C and the fact that $\tau \in \mathsf{C}_{\mathsf{power}}(\mathcal{P})$, it is a fetch event $(\mathsf{fetch}, \mathsf{tid}, i, \mathsf{instr})$. Choose the case based on the kind of instr. By NF-A and NF-B, all events related to the instruction (tid, i) constitute prefixes of $\overline{\tau^s_\mathsf{h}}$, $\mathsf{h} \in \mathsf{HEAD}$. The requirement $i = \mathsf{instr\text{-}count}(\mathsf{tid})$ holds by CMPL-R. The requirements like $\mathsf{h}_1 \leq \mathsf{h}_2 \leq \mathsf{h}_3$ in the load case naturally follow from the fact that $\tau \in \mathsf{C}_{\mathsf{power}}(\mathcal{P})$.

Assume $\tau^s_\mathsf{h} = \tau_\mathsf{h}$ for all $\mathsf{h} \in \mathsf{HEAD}$. Then $\tau^s_\mathsf{h} \in F_M$ by choice of $\mathsf{mem\text{-}value}_g$ and $\mathsf{last\text{-}key}_g$ in s_{M1} and CMPL-L, CMPL-M.

□

Lemma 18. *$\{\tau \in \mathsf{C}_{power}(\mathcal{P}) \mid \tau$ is in normal form of degree $n\} \subseteq \mathcal{L}(M(\mathcal{P})) \subseteq \mathsf{C}_{power}(\mathcal{P})$.*

Proof. Corollary of Lemmas 16 and 17. □

5.2 Checking Cyclicity of the Happens-Before Relation

We call a happens-before cycle *beautiful*, if it has the following form:

$$(\mathsf{tid}_1, i_1, \mathsf{instr}_1) \rightarrow_{po}{}^* (\mathsf{tid}_1, i'_1, \mathsf{instr}'_1) \rightarrow_{hop} \dots \rightarrow_{hop} (\mathsf{tid}_n, i_n, \mathsf{instr}_n) \rightarrow_{po}{}^* (\mathsf{tid}_n, i'_n, \mathsf{instr}'_n) \rightarrow_{hop} (\mathsf{tid}_1, i_1, \mathsf{instr}_1).$$

Here, $\rightarrow_{hop} := (\rightarrow_{co} \cup \rightarrow_{src} \cup \rightarrow_{cf})$ and $\mathsf{tid}_k \neq \mathsf{tid}_l$ for $k \neq l$. We call $\theta := \mathsf{tid}_1 \ldots \mathsf{tid}_n$ the *profile* of the cycle.

Example 4. The happens-before cycle shown in Figure 2 is beautiful.

Lemma 19 ([8]). *A computation $\tau \in \mathsf{C}_{\mathsf{power}}(\mathcal{P})$ has a happens-before cycle iff it has a beautiful happens-before cycle.*

Given a cycle profile θ, we define the automaton $M'(\mathcal{P}, \theta)$ as a modification of $M(\mathcal{P})$ that marks one event in each thread $\mathsf{tid}_j \in \theta$ by enter (identifying $(\mathsf{tid}_j, i_j, *)$) and a later (or the same) event by leave (identifying $(\mathsf{tid}_j, i'_j, *)$, $i_j \leq i'_j$). Note that $M(\mathcal{P})$ generates the events in program order, which ensures $(\mathsf{tid}_j, i_j, *) \rightarrow_{po}{}^* (\mathsf{tid}_j, i'_j, *)$. Technically, $M'(\mathcal{P}, \theta)$ introduces the following changes:

- The alphabet is $\mathsf{E}' := \mathsf{E} \times 2^{\{\mathsf{enter}, \mathsf{leave}\}}$.
- The automaton generates only load and prop events, as only they are relevant for cycle detection.
- The prop events include k component of the corresponding commit event.

To check $(\mathsf{tid}_j, i'_j, *) \rightarrow_{hop} (\mathsf{tid}_{j+1}, i_{j+1}, *)$, we use an intersection with a regular language $H^{\mathsf{tid}_j, \mathsf{tid}_{j+1}}$. The language $H^{\mathsf{tid}_1, \mathsf{tid}_2}$ includes a computation τ iff one or more of the following conditions hold:

H-ST $(\mathsf{e}_1, m_1), (\mathsf{e}_2, m_2) \in \tau$, $\mathsf{leave} \in m_1$, $\mathsf{enter} \in m_2$, $\mathsf{e}_1 = (\mathsf{prop}, \mathsf{tid}_1, \mathsf{tid}_1, \mathsf{k}_1, \mathsf{a})$, $\mathsf{e}_2 = (\mathsf{prop}, \mathsf{tid}_2, \mathsf{tid}_2, \mathsf{k}_2, \mathsf{a})$, and $\mathsf{k}_1 < \mathsf{k}_2$.

H-SRC $\tau = \tau_1 \cdot (\mathsf{e}_1, m_1) \cdot \tau_2 \cdot (\mathsf{e}_2, m_2) \cdot \tau_3$, $\mathsf{leave} \in m_1$, $\mathsf{enter} \in m_2$, $\mathsf{e}_1 = (\mathsf{prop}, \mathsf{tid}_2, \mathsf{tid}_1, \mathsf{a})$, $\mathsf{e}_2 = (\mathsf{load}, \mathsf{tid}_2, \mathsf{a})$, τ_2 does not contain events $(\mathsf{prop}, \mathsf{tid}_2, *, \mathsf{a})$.

H-CF1 $\tau = \tau_1 \cdot (\mathsf{e}_3, m_3) \cdot \tau_2 \cdot (\mathsf{e}_2, m_2) \cdot \tau_3$, $\mathsf{leave} \in m_2$, $\mathsf{e}_3 = (\mathsf{prop}, \mathsf{tid}_1, \mathsf{tid}_3, \mathsf{k}_3, \mathsf{a})$, $\mathsf{e}_2 = (\mathsf{load}, \mathsf{tid}_1, \mathsf{a})$, τ_2 does not contain events $(\mathsf{prop}, \mathsf{tid}_1, *, *, \mathsf{a})$, $(\mathsf{e}_3, m_3) \in \tau_1 \cdot \tau_2 \cdot \tau_3$, $m_3 \in \mathsf{enter}$, $\mathsf{e}_3 = (\mathsf{prop}, \mathsf{tid}_2, \mathsf{tid}_2, \mathsf{k}_2)$, $\mathsf{k}_3 < \mathsf{k}_2$.

H-CF2 $(\mathsf{e}_1, m_1), (\mathsf{e}_2, m_2) \in \tau$, $\mathsf{enter} \in m_1$, $\mathsf{leave} \in m_2$, $\mathsf{e}_1 = (\mathsf{load}, \mathsf{tid}_1, \mathsf{a})$, $\mathsf{e}_2 = (\mathsf{prop}, \mathsf{tid}_2, \mathsf{tid}_2, \mathsf{k}_2, \mathsf{a})$ and there is no $(\mathsf{e}_3, m_3) \in \tau$ with $\mathsf{e}_3 = (\mathsf{prop}, \mathsf{tid}_3, \mathsf{tid}_3, \mathsf{k}_3, \mathsf{a})$ with $\mathsf{k}_3 < \mathsf{k}_2$.

Lemma 20. *Program $\mathcal{P}$ has a beautiful cycle with profile $\theta = \mathsf{tid}_1 \ldots \mathsf{tid}_n$ iff*

$$M'(\mathcal{P}, \theta) \cap H^{\mathsf{tid}_1, \mathsf{tid}_2} \cap \ldots \cap H^{\mathsf{tid}_n, \mathsf{tid}_1} \neq \emptyset.$$

Note that $M'(\mathcal{P}, \theta)$ is infinite-state. To ensure $M'(\mathcal{P}, \theta)$ has finitely many states, we note that the instruction indices are irrelevant for the detection of happens-before cycles ($\mathsf{instr\text{-}count}$ can be dropped), and that the number of different coherence keys that must be stored in the state at any moment is polynomial in the size of $\mathcal{P}$. Indeed, the $\mathsf{last\text{-}key}$ and $\mathsf{last\text{-}key}_g$ components of the state each store at most $|\mathsf{ADDR}| \cdot |\mathcal{P}| \cdot n$ different coherence keys. Each modification of the $\mathsf{last\text{-}key}$ component of the state can be extended by a normalization step that would turn coherence keys to consecutive natural numbers starting from zero.

The normalization step must preserve the less-than relation on the keys. In order for the detection of happens-before cycles to work correctly, the automaton has to remember the coherence keys of marked store events: they must be preserved during normalization. Altogether, this results into $O(|\mathsf{ADDR}| \cdot |\mathcal{P}|^2 \cdot n)$ different keys, which is polynomial in the size of $\mathcal{P}$.

Theorem 2. *Robustness against Power is* PSPACE*-complete.*

Proof. By Theorem 1, Lemma 19, and Lemma 20, a program is non-robust iff the equation from Lemma 20 holds for some θ. In order to check robustness, we enumerate all profiles θ and check the equation from Lemma 20. The enumeration can be done in PSPACE. By construction and Lemma 14, the size of the intersection automaton is exponential in the size of the program. By Lemma 15, language emptiness for it can be checked in PSPACE in the size of the program, which gives us the upper bound.

The PSPACE lower bound follows from PSPACE-hardness of SC state reachability. One can reduce reachability to robustness by inserting an artificial happens-before cycle in the target state. □

Acknowledgements. The authors thank Parosh Aziz Abdulla, Jade Alglave, Mohamed Faouzi Atig, Ahmed Bouajjani, and Carl Leonardsson for helpful discussions on the Power memory model and the anonymous reviewers for suggestions. The first author was granted by the Competence Center High Performance Computing and Visualization (CC-HPC) of the Fraunhofer Institute for Industrial Mathematics (ITWM). The work was partially supported by the DFG project R2M2: Robustness against Relaxed Memory Models.

References

1. J. Alglave, October 2013. Personal communication.
2. J. Alglave and L. Maranget. Stability in weak memory models. In *CAV*, volume 6806 of *LNCS*, pages 50–66. Springer, 2011.
3. J. Alglave, L. Maranget, and M. Tautschnig. Herding cats. *CoRR*, abs/1308.6810, 2013.
4. A. Bouajjani, E. Derevenetc, and R. Meyer. Checking and enforcing robustness against TSO. In *ESOP*, volume 7792 of *LNCS*, pages 533–553. Springer, 2013.
5. A. Bouajjani, R. Meyer, and E. Möhlmann. Deciding robustness against Total Store Ordering. In *ICALP*, volume 6756 of *LNCS*, pages 428–440. Springer, 2011.
6. S. Burckhardt and M. Musuvathi. Effective program verification for relaxed memory models. In *CAV*, volume 5123 of *LNCS*, pages 107–120. Springer, 2008.
7. J. Burnim, C. Stergiou, and K. Sen. Sound and complete monitoring of sequential consistency for relaxed memory models. In *TACAS*, volume 6605 of *LNCS*, pages 11–25. Springer, 2011.
8. G. Calin, E. Derevenetc, R. Majumdar, and R. Meyer. A theory of partitioned global address spaces. In *FSTTCS*, volume 24 of *LIPIcs*, pages 127–139, 2013.
9. L. Lamport. Time, clocks, and the ordering of events in a distributed system. *CACM*, 21(7):558–565, 1978.

10. L. Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Transactions on Computers*, 28(9):690–691, 1979.
11. S. Mador-Haim, L. Maranget, S. Sarkar, K. Memarian, J. Alglave, S. Owens, R. Alur, M. M. K. Martin, P. Sewell, and D. Williams. An axiomatic memory model for POWER multiprocessors. In *CAV*, volume 7358 of *LNCS*, pages 495–512. Springer, 2012.
12. L. Maranget, S. Sarkar, and P. Sewell. A tutorial introduction to the ARM and POWER relaxed memory models. `https://www.cl.cam.ac.uk/~pes20/ppc-supplemental/test7.pdf`. Draft.
13. S. Owens, S. Sarkar, and P. Sewell. A better x86 memory model: x86-TSO (extended version). Technical Report CL-TR-745, University of Cambridge, 2009.
14. S. Sarkar, P. Sewell, J. Alglave, L. Maranget, and D. Williams. Understanding POWER multiprocessors. In *PLDI*, pages 175–186. ACM, 2011.
15. D. Shasha and M. Snir. Efficient and correct execution of parallel programs that share memory. *ACM TOPLAS*, 10(2):282–312, 1988.